



USING SHADOWS TO DETECT TARGETS
IN SYNTHETIC APERTURE RADAR IMAGERY

THESIS

Brian Donnell, Captain, USAF

AFIT/GE/ENG/09-12

DEPARTMENT OF THE AIR FORCE
AIR UNIVERSITY

AIR FORCE INSTITUTE OF TECHNOLOGY

Wright-Patterson Air Force Base, Ohio

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

The views expressed in this thesis are those of the author and do not reflect the official policy or position of the United States Air Force, Department of Defense, or the United States Government.

AFIT/GE/ENG/09-12

USING SHADOWS TO DETECT TARGETS
IN SYNTHETIC APERTURE RADAR IMAGERY

THESIS

Presented to the Faculty
Department of Electrical and Computer Engineering
Graduate School of Engineering and Management
Air Force Institute of Technology
Air University
Air Education and Training Command
In Partial Fulfillment of the Requirements for the
Degree of Master of Science in Electrical Engineering

Brian Donnell, B.S.E.E.
Captain, USAF

March 2009

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

USING SHADOWS TO DETECT TARGETS
IN SYNTHETIC APERTURE RADAR IMAGERY

Brian Donnell, B.S.E.E.
Captain, USAF

Approved:

/signed/	13 Mar 2009
Michael J. Mendenhall, Maj, PhD (Chairman)	date
/signed/	13 Mar 2009
Michael A. Saville, Maj, PhD (Member)	date
/signed/	13 Mar 2009
Dr. Richard K. Martin (Member)	date

Abstract

Synthetic Aperture Radar (SAR) can generate high resolution imagery of remote scenes by combining the phase information of multiple radar pulses along a given path. SAR based Intelligence, Surveillance, and Reconnaissance (ISR) has the advantage over optical ISR that it can provide usable imagery in adverse weather or nighttime conditions. Certain radar frequencies can even result in foliage or limited soil penetration, enabling imagery to be created of objects of interest that would otherwise be hidden from optical surveillance systems.

This thesis demonstrates the capability of locating stationary targets of interest based on the locations of their shadows and the characteristics of pixel intensity distributions within the SAR imagery. Shadows, in SAR imagery, represent the absence of a detectable signal reflection due to the physical obstruction of the transmitted radar energy. An object's shadow indicates its true geospatial location. This thesis demonstrates target detection based on shadow location using three types of target vehicles, each located in urban and rural clutter scenes, from the publicly available Moving and Stationary Target Acquisition and Recognition (MSTAR) data set.

The proposed distribution characterization method for detecting shadows demonstrates the capability of isolating distinct regions within SAR imagery and using the junctions between shadow and non-shadow regions to locate individual shadow-casting objects. Targets of interest are then located within that collection of objects with an average detection accuracy rate of 93%. The shadow-based target detection algorithm results in a lower false alarm rate compared to previous research conducted with the same data set, with 71% fewer false alarms for the same clutter region. Utilizing the absence of signal, in conjunction with surrounding signal reflections, provides accurate stationary target detection. This capability could greatly assist in track initialization or the location of otherwise obscured targets of interest.

Acknowledgements

Receiving a graduate degree was something I never expected to do when I completed my Bachelor's in electrical engineering. The process has been long, tiring, and more difficult than I expected. It reminded me of all the things I both like and dislike about any education program. But ultimately, it was a rewarding experience.

First, and most important, I thank God. He not only provided the opportunities that led me to this point, but granted me the wisdom to see the correct path before me and the strength to follow it. Having served on base Honor Guard, I have helped bury more fellow Airmen and soldiers than I care to recount. Yet in these troubled times of conflict I have been graced with the assignment of completing a higher education program. I believe the 16th century politician John Bradford said it best with the line, "There but for the grace of God, go I."

Second, I would like to thank Lieutenant Colonel Dave Chelen. Without his faith in my abilities I would (quite literally) never have been accepted to AFIT. His actions relating to my enrollment at AFIT will serve as an example of an outstanding commander for the remainder of my career, and beyond.

Finally, I would like to thank my family. Though they may not have had an understanding of the topic of this thesis, they still supported my efforts and provided emotional support whenever possible.

Brian Donnell

Table of Contents

	Page
Abstract	iv
Acknowledgements	v
Table of Contents	vi
List of Figures	viii
List of Tables	x
List of Abbreviations	xi
 I. Introduction	 1-1
1.1 Research Motivation	1-1
1.2 Problem Statement	1-2
1.3 Research Scope	1-3
1.4 Organization	1-3
 II. Static and Dynamic Target Detection in Synthetic Aperture Radar (SAR)	 2-1
2.1 Standard Radar	2-1
2.1.1 Standard Range Resolution	2-2
2.1.2 Standard Cross Range Resolution	2-4
2.2 Extending the Radar Equations to SAR	2-5
2.2.1 The Projection-Slice Theorem	2-6
2.2.2 SAR Range and Cross Range Resolutions	2-8
2.3 SAR Image Reconstruction	2-10
2.3.1 Polar Reformatting Algorithm	2-11
2.3.2 Convolution Back-Projection Algorithm	2-11
2.4 Motion Effects in SAR Imagery	2-13
2.5 Change Detection in SAR Imagery	2-16
2.5.1 Interferometric SAR (IF-SAR)	2-16
2.5.2 Displaced Phase Center Algorithm (DPCA)	2-17
2.6 Shadow Detection	2-19
2.6.1 Previous Work Exploiting Shadow Information in SAR Data	2-20
2.6.2 Previous Work in Shadow Detection in SAR Data	2-21
2.6.3 Scene Statistics	2-22
2.7 Summary	2-24

	Page
III. Research Methodology	3-1
3.1 Shadow Detection	3-1
3.1.1 Scene Characterization For Distribution Statistics	3-2
3.1.2 Generate Shadow Mask	3-4
3.1.3 Low-pass Speckle Filter	3-4
3.1.4 Shadow-to-Non-Shadow Detection	3-5
3.1.5 Group Transition Points	3-8
3.2 Filtering For Targets of Interest	3-10
3.2.1 Filter By Group Size and Combine Overlapping Targets	3-10
3.2.2 Filter By Distribution Properties	3-13
3.3 Determining the Optimum Filter Parameters	3-14
3.4 Summary	3-18
IV. Experiment Methodology	4-1
4.1 Data Set	4-1
4.1.1 Data Implementation	4-2
4.1.2 Data Selection	4-3
4.2 Locating Targets From Shadows	4-5
4.2.1 Scene Characterization for Distribution Statistics	4-7
4.2.2 Generate and Low-pass Speckle Filter the Shadow Mask	4-9
4.2.3 Shadow-to-Non-Shadow Detection and Group Tran- sition Points	4-13
4.3 Filter For Targets of Interest	4-15
4.3.1 Filter by Group Size	4-16
4.3.2 Combine Overlapping Targets	4-18
4.3.3 Extract Target Intensity Statistics	4-20
4.4 Training the Filter	4-22
4.4.1 Minimization Function	4-22
4.4.2 Simulated Annealing Results	4-23
4.4.3 K -Fold Cross Validation	4-25
4.5 Mixed Targets	4-27
4.6 Summary	4-29
V. Conclusions	5-1
5.1 Research Contributions	5-1
5.2 Future Research	5-2
Bibliography	BIB-1

List of Figures

Figure		Page
2.1.	Diagram of an Uncompressed Radar Range Resolution.	2-3
2.2.	Diagram of Strip-Mapping SAR.	2-6
2.3.	Diagram of Spotlight SAR.	2-7
2.4.	Diagram of the Projection-Slice Theorem.	2-8
2.5.	Spatial Frequency Support Region for a SAR Collection.	2-9
2.6.	Interpolating Polar Data Points to a Rectangular Grid.	2-12
2.7.	SAR Generated Images of Stationary and Moving Point Targets.	2-14
2.8.	Examples of Target Motion Effects in SAR Imagery.	2-15
2.9.	Displaced Phase Center Algorithm Implementation.	2-18
2.10.	Shadows Revealed by DPCA.	2-20
3.1.	A Block Diagram of the Target Detection Algorithm.	3-2
3.2.	Magnitude Response of a Typical Butterworth Filter.	3-6
3.3.	Image Orientation With Respect to Sensor Platform.	3-7
3.4.	Grouping Transition Pixels Process.	3-9
3.5.	Examples of Depression Angle Affecting the Amount of Self-Shadowing.	3-11
3.6.	Single Target With Two Shadow Transition Lines.	3-12
3.7.	Gamma Distributions With Constant Scale and Varying Shape.	3-14
3.8.	Plots of P_S For Different Values of Detections and False Alarms.	3-17
3.9.	Simulated Annealing Pseudo-code.	3-18
4.1.	Selected SAR Image Clutter Scenes.	4-4
4.2.	Examples of the Target Chips for Each Vehicle Type.	4-6
4.3.	Histograms of the Distributions of Selected Image Regions.	4-8
4.4.	Distribution Plots of Selected Regions of the Clutter Scene	4-9

Figure		Page
4.5.	Comparison of Probability Density Functions for Target and Clutter Shadows	4-10
4.6.	Unfiltered Shadow Mask of the Rural Clutter Scene	4-10
4.7.	One Dimensional Butterworth Low-pass Filter Responses for Varying Orders of the Filter.	4-11
4.8.	Butterworth Filtered Shadow Masks.	4-12
4.9.	Filtered Shadow Mask of a Target Chip	4-13
4.10.	Shadow-to-Non-Shadow Transitions in the Rural Scene for the T-72 Height Parameter	4-14
4.11.	Filtering Shadow Groups by Size.	4-17
4.12.	Combining Overlapping Targets in a Single Target Chip.	4-19
4.13.	Histogram Comparison of Known Targets and a False Alarm.	4-21
4.14.	BMP-2 P_S as a Function of Only the Pixel Percentages u and n	4-24
4.15.	Sparse Target Detection For All Three Types of Vehicles in Each Scene.	4-28

List of Tables

Table		Page
4.1.	Target Vehicle Dimensions	4-2
4.2.	Number of Transitions for Each Vehicle in Each Scene	4-15
4.3.	Number of Transition Groups for Each Vehicle in Each Scene	4-15
4.4.	Shadow Group Limits, in Pixels, for Each Vehicle	4-17
4.5.	Size-Filtered Shadow Group Target Detection Rates	4-18
4.6.	Target Detection Rates After Combining Overlapping Bounding Boxes	4-20
4.7.	Initial Simulated Annealing Results for the BMP-2 Target	4-24
4.8.	5-Fold Cross Validation Results for the T-72 Data	4-26
4.9.	5-Fold Cross Validation Results for the BMP-2 Data	4-26
4.10.	5-Fold Cross Validation Results for the BTR-70 Data	4-26
4.11.	Target Detection Results for Combined Target Scenes	4-29

List of Abbreviations

Abbreviation		Page
SAR	Synthetic Aperture Radar	1-1
DSP	Digital Signal Processing	2-1
LFM	Linear Frequency Modulation	2-3
CAT	Computerized Axial Tomography	2-6
CBP	Convolution Back-Projection	2-11
IF-SAR	Interferometric SAR	2-16
DPCA	Displaced Phase Center Algorithm	2-17
PDF	Probability Distribution Function	2-22
RCS	Radar Cross Section	2-22
MSTAR	Moving and Stationary Target Acquisition and Recognition	3-10
DARPA	Defense Research Advanced Projects Agency	4-1
AFRL	Air Force Research Laboratory	4-1
SNL	Sandia National Laboratory	4-1
GIS	Geographic Information System	5-3

USING SHADOWS TO DETECT TARGETS IN SYNTHETIC APERTURE RADAR IMAGERY

I. Introduction

Synthetic Aperture Radar (SAR) is a method of generating high resolution surveillance imagery at long ranges. Common uses for SAR include measuring terrain deformation and elevation, monitoring land use, and generating remote imagery of hostile regions. SAR sensors are mounted on both airborne and spaceborne platforms and have been in use for over fifty years as a reliable method of remote sensing.

Resolutions of SAR imagery can be on par with optical systems, but SAR provides a number of distinct advantages. SAR operates by transmitting electromagnetic energy and measuring the amount reflected by everything in the illuminated scene. Different frequencies will have different properties, and allow for specialized uses. For instance, lower frequencies can penetrate foliage and image features that would be optically obstructed by trees or other vegetation, which is extremely useful for search and rescue or law enforcement. Even lower frequencies have been shown to have limited ground penetration capabilities, providing imagery of shallow buried objects or utility piping. More common uses for SAR, that would prove problematic for an optical sensor, include imaging during adverse weather conditions or night operations. SAR can “see” through cloud cover, rain, or dust storms. While optical sensors are limited to the amount of light in a scene, SAR is an active sensor providing its own illumination. This means a SAR platform can generate the same quality imagery regardless of the time of day during the collection.

1.1 Research Motivation

As mentioned at the start of this chapter, SAR does provide high resolution imagery of remote scenes. However, this is typically restricted to stationary objects. One of the limitations of SAR is that, due to the principles allowing image construc-

tion, moving objects in a scene will appear blurred and/or displaced from their true locations in the final image.

Restricting SAR image analysis to stationary objects can provide valuable information (such as the terrain, building layout, construction, or vegetation in an area), but certain remote sensing objectives require the ability to locate or identify moving objects. For instance, search and rescue, traffic analysis, law enforcement, and military operations can all benefit from the ability to locate or identify moving objects.

Considerable research has been conducted in the realm of refocusing and correctly positioning moving objects in SAR imagery. However, the vast majority of papers focus on the phase information contained in the radar signal returned from the reflection of the moving object [12, 13, 21, 29, 30]. These methods allow for focusing and positioning of moving objects, but they require the motion parameters of the object to be known. If the parameters are unknown, images must be generated for the range of the potential values of those parameters. This can be an extremely time consuming and computationally expensive process.

1.2 Problem Statement

Since SAR is an active remote sensing method, transmitting electromagnetic energy to illuminate a scene, objects cast shadows as the transmitted energy is reflected off them and prevented from continuing. The previously mentioned problem of moving objects appearing blurred and displaced in SAR imagery is more fully explained in Section 2.4, but the fundamental cause has to do with the phase information of the reflected radar energy. Shadows, as the absence of reflected energy, have no phase information, and are therefore not displaced in the generated image.

This thesis attempts to locate objects in pre-generated SAR imagery based on the shadows within the image, and then identify particular targets of interest from the remainder of the located objects. First, it must be demonstrated that shadows

can be reliably isolated from non-shadows in SAR imagery. Second, a method of locating the objects that created the shadows must be developed. Finally, a method of differentiating and identifying particular targets of interest from among all of the shadow casting objects in the scene must be demonstrated. The final result will be the locations of desired targets of interest within a SAR image.

1.3 Research Scope

Based on the literature review conducted for Sections 2.6.1 and 2.6.2, this research focuses on locating stationary objects based on shadows. Given how little publicly available previous research has been conducted regarding shadows in SAR imagery, it must first be determined if shadows provide enough information for target location in the simplest of cases before attempting it in more complex scenarios. A combination of distribution analysis and image processing will be used to identify and locate objects in the imagery. Then, an optimization algorithm will assist in training a decision function to separate targets of interest from background clutter objects.

1.4 Organization

This thesis is organized into five chapters. Aside from this introduction chapter, the following four chapters each address a specific aspect of the research. Chapter II presents the radar and SAR principles utilized in the development of the algorithm presented in this thesis. Then, previous work regarding shadows in SAR imagery is detailed and each work's applicability to the problem at hand is explained.

Chapter III details the research methodology used in developing and demonstrating the target location algorithm. An overview of the algorithm is presented and then each step of the process is explained in detail.

Chapter IV follows each step in the target location algorithm again, presenting the results of each operation. Each stage of the shadow detector's construction is

illustrated through the first half of the algorithm. Finally, the impact on overall target detection accuracy is listed for each step of the target filter.

Chapter V contains a summary of the conclusions of this research and outlines a selection of potential future research topics.

II. Static and Dynamic Target Detection in Synthetic Aperture Radar (SAR)

Synthetic Aperture Radar (SAR) has been around for over fifty years. However, it is recent advances in computers and digital signal processing (DSP) that are allowing new implementations and applications of SAR data. This chapter presents the basic principles behind radar, then expands these principles to SAR image generation. With a basic understanding of how SAR images are generated, some of the problems created by moving targets within the collection scene are explained. The more popular DSP methods of detecting and separating moving targets from the static returns, to minimize these previously mentioned problems, are then discussed. Finally, previous work is detailed in using shadows, rather than returns, of targets to perform target detection in order to eliminate problems encountered due to moving targets.

2.1 *Standard Radar*

Radar, in its most basic form, is described in [27] as detecting the range of an object by transmitting an electromagnetic pulse and measuring how long it takes the echo of that pulse to return. More precisely, an electromagnetic pulse is transmitted by an antenna into space. This pulse travels at the speed of light for the medium it is transiting until it encounters an object or discontinuity in the spacial medium. At this point, some of the original energy is reflected while the remainder travels through or excites the impinged object, resulting in excitation and retransmission. This combined return again travels at the speed of light back through the medium. Some of the reflected and retransmitted energy will be detected by the receiving antenna and the time between transmit and receive can be used to calculate the distance to the object. If the speed of light for the transmitting medium is known, and assuming the same antenna is used for transmitting and receiving, this range is determined by

$$R = \frac{cT}{2} \quad (2.1)$$

where c is the speed of light through the transmission medium, T is the time between pulse transmission and reception, and R is the distance from the radar to the object.

2.1.1 Standard Range Resolution. Since radar operates by transmitting and receiving a finite length pulse, the resolution at which two separate targets can be determined in range is based on the effective transmission length of that pulse. For instance, if an unmodulated pulse hits one target, part of the energy will begin reflecting. If the leading edge of the same pulse impacts a second target, and the front of that reflection reaches the leading edge of the first target before the trailing edge of the pulse is fully past it, the combined reflection of the two will appear as a single target to the radar. This means the minimum range resolution for an unmodulated pulse is

$$\delta R = \frac{c\tau}{2} \quad (2.2)$$

where τ is the time required to transmit the pulse, also known as the pulse width. An example of a radar pulse passing two objects at the very edge of its range resolution limit is illustrated in Figure 2.1. It becomes apparent that smaller pulse widths mean increased range resolution. However, smaller pulses also mean that higher peak transmitting powers are required to maintain the same average power for a pulse. The power directly corresponds to the maximum operating range of the radar as the transmitted power must be high enough to allow a detectable level to return to the receiver after reflecting from a target. The most common method to balance both power limitations and effective pulse lengths is through pulse compression.

Pulse compression achieves shorter effective pulse lengths for the same actual pulse length by modulating the transmitted pulse waveform in either frequency or phase. Since the actual pulse length is the same, the power requirements are com-




Figure 2.1: In this example, the two objects are too close together to be discerned as separate targets for the given pulse width. The reflections of each object slightly overlap and the radar will only “see” one return.

parable to that of an uncompressed or unmodulated pulse. Radar receivers typically process the received signal through a matched filter matched to the transmitted pulse. This effectively implements a convolution operation across the pulse with a mirror of itself. This means for a positive return the detectable energy level of a compressed pulse is over a shorter time span, meaning an increased range resolution without a loss of other performance parameters.

Many forms of pulse compression are described in [16], but the most common form in use is Linear Frequency Modulation (LFM), also known as a chirp signal. Because it is the most common it is the one used throughout this thesis. The unmodulated pulse is defined as the rectangular function, $u(t)$, with a width equal to the duration of the pulse transmission. However, the LFM pulse is defined by Jakowatz, et al. in [11] as

$$s(t) = \text{Re}\{\exp[j(\omega_0 t + \alpha t^2)]\} \quad (2.3)$$

over the interval $-\tau_c/2 \leq t \leq \tau_c/2$ where τ_c is the length of the pulse, ω_0 is the base frequency of the pulse, $\omega_0 t$ is the linear phase term, α is the frequency modulation rate, and αt^2 is the quadratic phase term. Jakowatz proves that the range resolution for an LFM pulse is determined by its frequency bandwidth B as

$$\delta R = \frac{c}{2B}. \quad (2.4)$$

This form of the range resolution equation provides better resolution for the same pulse width given that the bandwidth is large enough.

2.1.2 Standard Cross Range Resolution. Until now, the cross range aspect of the standard radar system has been ignored. If the antenna used by the radar is an ideal isotropic radiator there would be no cross range resolution possible as the transmitted pulse would radiate in all directions equally. To combat this effect, the radar antenna is designed with the primary purpose of the radar in mind. For instance, early warning radar will typically use a wide beamwidth to scan as large an area as possible. On the other hand, search radar will have a very narrow beamwidth and rotate the antenna to cover the desired area. If a range detection occurs, the cross range resolution is only as good as the antenna beamwidth. The angular width of a beam is directly proportional to the operating wavelength and antenna aperture size such that

$$\beta = \frac{\lambda}{D} \quad (2.5)$$

where β is the angular beamwidth in radians, λ is the operating frequency wavelength, and D is the aperture width. This means the cross range resolution for a given distance R is

$$\delta R_{cr} = \frac{R\lambda}{D} \quad (2.6)$$

For a sub-meter cross range resolution, at standard radar operating frequencies and any appreciable distance, the required antenna aperture width exceeds 1 km. For example, at the extremely close range of 1 km operating at X-band with a wavelength of 3 cm, the required aperture width is

$$D = \frac{R\lambda}{\delta R_{cr}} = \frac{(1000m)(0.03m)}{(1m)} = 30m \quad (2.7)$$

While 30 meters, or 100 feet, is not exceedingly large for ground-based radar systems, it is far too large to realistically mount on an aircraft. Furthermore, increasing the operating range increases the required aperture size.

2.2 Extending the Radar Equations to SAR

A solution to the cross range resolution problem was first devised by Carl Wiley in the 1950s, known as Doppler beam sharpening [11]. Doppler beam sharpening is now known as strip-mapping SAR, where the radar antenna is mounted orthogonal to the aircraft flight path and pulses are transmitted and collected as it passes the target scene, as shown in Figure 2.2. Each of these returns are combined and processed to form a return from a much larger “synthetic aperture.” For this strip-mapping synthetic aperture, the largest possible physical antenna beamwidth is desired, in order to maximize the number of returns for a specific point in space as the radar moves. This is opposite to the standard radar convention. An antenna’s beamwidth is typically inversely proportional to its physical aperture size. This means where standard radar desires a large physical antenna, as described in Section 2.1.2, a synthetic aperture radar works best with the smallest physical antenna size possible. However, the required sampling rate to prevent aliasing provides a lower bound for how small the antenna can be, as samples must be acquired at least every half-width of the physical antenna.

The second type of SAR collection is known as spotlight SAR. In strip-mapping SAR the aircraft must fly the length of the target scene plus the length of the antenna




Figure 2.2: In strip-mapping SAR, radar pulses are transmitted orthogonal to the aircraft’s flight path, sampling a “strip” of ground.

footprint to insure the best cross range resolution possible. Spotlight SAR flies in a circular pattern over the target scene, as shown in Figure 2.3. The entire scene is always illuminated by the antenna, resulting in shorter flight paths for the same cross range resolution. This circular collection closely parallels the collection method used in medical computerized axial tomography (CAT) [11].

2.2.1 The Projection-Slice Theorem. Medical X-ray tomography relies on the relationship that the one-dimensional Fourier transform of a projection function, $p_{\theta}(u)$, is equal to the two-dimensional Fourier transform of an image, $g(x, y)$, evaluated along a radial line of the same θ , where θ is the angle of the radial line off the X axis, and u is the radius to any point on that line. This relationship is known as




Figure 2.3: In spotlight SAR, the aircraft flies in a circle around the scene to be imaged and every radar pulse samples the entire scene.

the Projection-Slice Theorem, and is illustrated in Figure 2.4. The Projection-Slice Theorem is described in [11] and written as

$$G(U\cos\theta, U\sin\theta) = P_\theta(U) \quad (2.8)$$

where $P_\theta(U)$ is the Fourier transform of the projection function $p_\theta(u)$

$$P_\theta(U) = \int_{-\infty}^{\infty} p_\theta(u) e^{-juU} du. \quad (2.9)$$

$G(U\cos\theta, U\sin\theta)$ is the two-dimensional Fourier transform of the image $g(x, y)$

$$G(X, Y) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} g(x, y) e^{-j(xX+yY)} dx dy \quad (2.10)$$




Figure 2.4: The Projection-Slice Theorem establishes a relationship between a two-dimensional image function and the radial projections of that function. This diagram is a common representation of the Projection-Slice Theorem found in texts.

where $x = u \cos \theta$ and $y = u \sin \theta$.

In X-ray tomography, the image, $g(x, y)$, is an attenuation coefficient profile, while in SAR it becomes a reflectivity map of the target scene.

2.2.2 SAR Range and Cross Range Resolutions. The Projection-Slice Theorem provides a basis for discussing image resolution in SAR. If the radar returns from a SAR collection are mapped to Fourier space, using the scene center as the reference point, the image resolution can be easily defined. The frequency modulation of each pulse means each individual frequency response corresponds to a unique range from the radar to a point in the scene. Each pulse is mapped to a radial line in the Fourier domain based on θ , the angle between the direction of the pulse's collection and a reference axis in the scene. The frequency of each data point, in radians, then represents




Figure 2.5: Each frequency sample from each pulse's return can be mapped to a unique point in the Fourier domain (spatial frequency). The result can be used to identify the extent and resolution of the image that can be created from the data.

the radius from that data point to the center of the domain. Mapping the collected data creates a segment of an annulus in the Fourier domain offset by $\frac{4\pi}{\lambda}$ where λ is the center frequency of the radar pulse, as shown in Figure 2.5. The range dimension, or radii of the annulus, spans the frequency bandwidth of the pulse

$$\Delta Y = \frac{2}{c}(2\pi B) \quad (2.11)$$

Taking the inverse Fourier transform of the radar returns then results in an image range resolution of

$$\delta y = \frac{2\pi}{\Delta Y} = \frac{c}{2B} \quad (2.12)$$

which is the same range resolution as the standard radar using an LFM pulse, as presented in Equation 2.4.

Similarly, the cross-range data spans the set of angles collected, $\Delta\theta$. The geometry of the annulus in polar coordinates means the length of the constant frequency arcs from one end of the annular segment to the other changes as the frequency changes. However, radar uses extremely high frequencies, meaning these changes in length are negligible for small angular spans. For simplicity, the center frequency, or offset of the annulus, is used as the ideal point to determine the span length resulting in

$$\Delta X = 2 \frac{4\pi}{\lambda} \sin\left(\frac{\Delta\theta}{2}\right). \quad (2.13)$$

In order to minimize error from platform motion and phase drift, the angular range used to generate an image is typically very small, meaning Equation 2.13 can be simplified using the small angle approximation as

$$\Delta X \approx \frac{4\pi}{\lambda} \Delta\theta. \quad (2.14)$$

The inverse Fourier transform of this term is

$$\delta x = \frac{2\pi}{\Delta X} = \frac{\lambda}{2\Delta\theta}. \quad (2.15)$$

Two things are immediately apparent when this result in Equation 2.15 is compared to Equation 2.6. The first is a lack of dependence on range to the target. The second is a complete separation from the size of the antenna itself. In this regard, the antenna size defines both the required sampling rate *and* the size of the scene imaged, but *does not* influence the cross range resolution.

2.3 SAR Image Reconstruction

Sections 2.2.1 and 2.2.2 described SAR pulse returns as spatial frequency projections of the target scene's reflectivity. In order for a human to view and process the spatial information, the data must be transformed into imagery. There are many dif-

ferent methods available for transforming SAR data into recognizable images, such as time domain correlation, range stacking, and the previously mentioned Doppler beam sharpening. Two of the more commonly used methods are presented here. The Polar Reformatting Algorithm is typically considered a relatively fast, if slightly inaccurate, method of image generation; while the Convolution Back-Projection Algorithm is at the other end of the spectrum as one of the most computationally intensive, but more accurate, routines.

2.3.1 Polar Reformatting Algorithm. One of the earliest image reconstruction techniques used for tomographic imaging is the polar reformatting algorithm [28]. Thanks to advances in the speed of modern computing it has become one of the faster, if less accurate, methods of SAR image reconstruction. Assuming the scene has been sampled at a high enough rate to prevent aliasing, the collections are mapped to the Fourier domain exactly as in Section 2.2.2. The data points in the annulus are then fitted to a uniform rectangular grid through interpolation and a two-dimensional inverse Fourier transform performed on the result, generating the final image of the scene. The inaccuracy of this method is a result of the original polar formatting, requiring a two-dimensional interpolation to create an evenly spaced rectangular grid of points that can be inverse Fourier transformed, as shown in Figure 2.6. Data points closer to the origin are more densely packed, but data points at higher spatial frequencies are further spread, resulting in less accurate interpolation.

2.3.2 Convolution Back-Projection Algorithm. The convolution back-projection (CBP) algorithm is a more direct implementation of the projection-slice theorem. As explained in [11], the connection is made by first expressing the two-dimensional Fourier transform of the image one wishes to construct in polar coordinates as

$$g(\rho \cos \phi, \rho \sin \phi) = \frac{1}{4\pi^2} \int_{-\pi/2}^{\pi/2} d\theta \int_{-\infty}^{\infty} G(r \cos \theta, r \sin \theta) |r| e^{jr \rho \cos(\phi - \theta)} dr \quad (2.16)$$




Figure 2.6: The polar reformatting algorithm requires interpolation in both rectangular spatial frequency dimensions. This can cause some inaccuracies in the resulting image because some of the intersection points in the grid are not near the measured data points.

where r is the radius from the origin to any point along the projection. When combining the results in Equation 2.16 with the projection-slice theorem, Equation 2.8 becomes

$$g(\rho \cos \phi, \rho \sin \phi) = \frac{1}{4\pi^2} \int_{-\pi/2}^{\pi/2} d\theta \int_{-\infty}^{\infty} P_{\theta}(r) |r| e^{jr\rho \cos(\phi - \theta)} dr. \quad (2.17)$$

The convolution portion of the algorithm is derived from the product of $P_{\theta}(r)$ with $|r|$ in the Fourier domain. The inverse Fourier transform of $|r|$ is the filtering kernel, $h(\rho)$. A multiplication in the Fourier domain is equal to a convolution in the image domain, and as such, Equation 2.17 can be rewritten as

$$g(\rho \cos \phi, \rho \sin \phi) = \frac{1}{4\pi^2} \int_{-\pi/2}^{\pi/2} Q(\rho \cos(\phi - \theta)) d\theta \quad (2.18)$$

where Q is a combination of the projection function, $p_\theta(u)$, convolved with the filtering function, $h(\rho)$. Converting Equation 2.18 back to Cartesian coordinates produces

$$g(x, y) = \frac{1}{4\pi^2} \int_{-\pi/2}^{\pi/2} Q(x\cos\theta + y\sin\theta) d\theta \quad (2.19)$$

which is the integral of the filtered projection functions back-projected along the line of their respective collections.

The CBP algorithm is implemented by creating a grid of coordinates in the image space and back-projecting each filtered projection function to these grid points by way of Equation 2.19. Since the grid points projected onto the projection function will not align perfectly with the sampled points, a one-dimensional interpolation is required. However, this interpolation is only in one dimension, as opposed to the two-dimensional interpolation required by the polar reformatting algorithm, and as such generates a more accurate image. Unfortunately, the many individual interpolations and back-projections for each sample angle greatly increases the number of operations, and hence time to generate an image, as compared to the polar reformatting algorithm.

2.4 *Motion Effects in SAR Imagery*

So far, in the explanation of how SAR imagery is formed, it has been assumed everything in the target scene remains stationary over the collection period. The two presented methods of generating SAR imagery are considerably different, yet affected similarly by motion in the scene during the collection. We will discuss how the reflection mapping used in the projection-slice theorem will change with each sampling if there is a moving object in the scene. It should be noted that the Projection-Slice Theorem assumes the mapping remains constant throughout the capturing process. As such, moving objects cause the Projection-Slice Theorem to break down because the image no longer remains constant.

To understand how a moving object corrupts a SAR image, it is beneficial to view the CBP algorithm as a summation of all the phase contributions of all objects

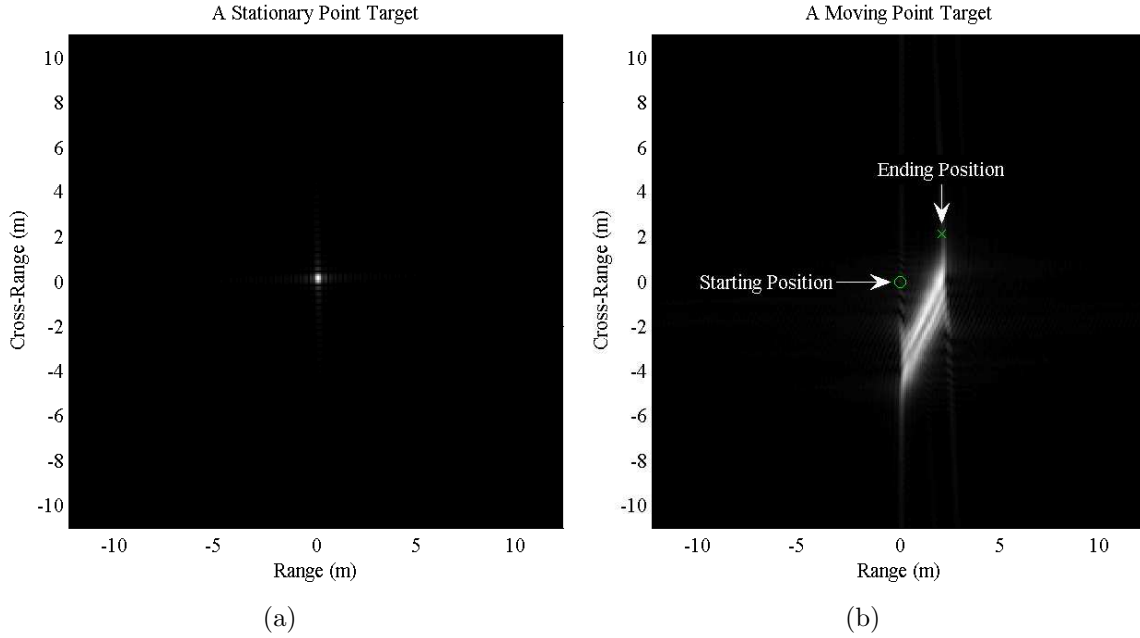


Figure 2.7: (a) The point target in this simulation is stationary and the image appears in the actual target location. (b) This point target is moving as the data is collected. This causes the image of the target to smear and displace.

in the scene across all the pulses. Stationary objects will coherently sum along the correct frequency wavefronts generating their representation in the image at their true location. If an object is moving, its phase contributions will no longer coherently sum in a meaningful fashion. This is because each “slice” used in the Projection-Slice Theorem is captured at a different point in time. If an object is moving spatially, the image represented by the slices will be constantly changing to represent to the location of the object at the specific point in time the slice was captured. Since the Projection-Slice Theorem relies on all the slices representing a single, constant, image the resulting generated image will be deformed. A comparison of images generated of stationary and moving point targets is shown in Figure 2.7.

Exactly how the image of the moving object is deformed is proportional to the radial velocity of the object relative to the radar. This is easier to visualize in strip-mapping SAR, but the principles are identical for spotlight SAR. For instance, if an object is moving with constant velocity normal to the radar’s motion (velocity purely

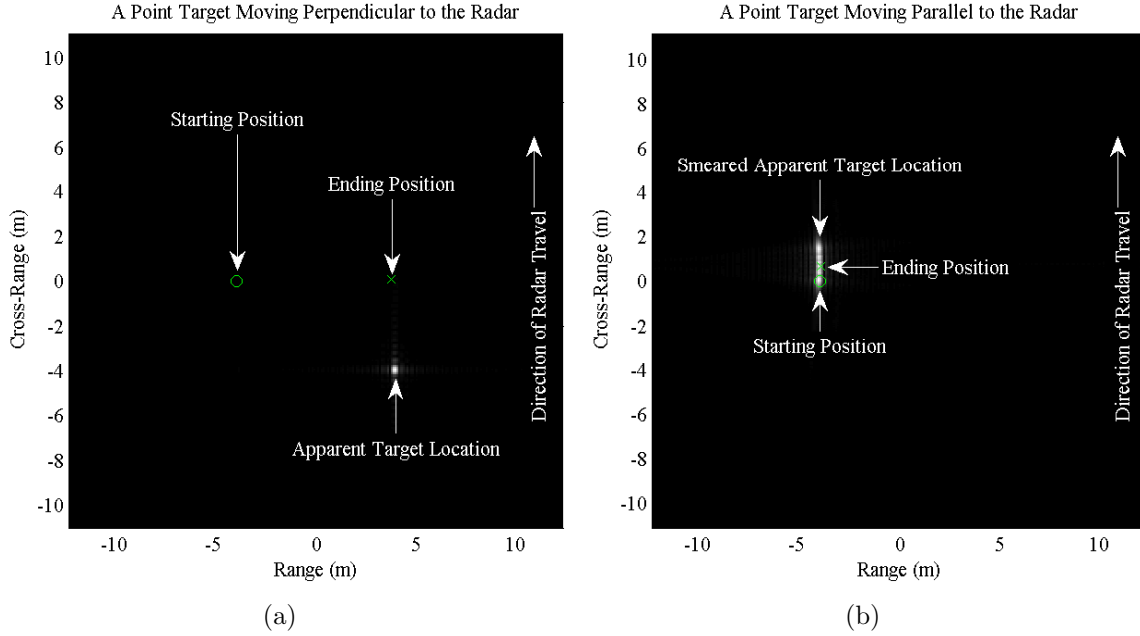


Figure 2.8: (a) The point target in this simulation is moving perpendicular to the radar’s motion and appears focused, but displaced from the path of motion. (b) This point target is moving parallel to the radar’s motion and appears smeared and unfocused, but the smear is aligned with the path of the target’s motion.

in the range direction), there will be a point where all the frequency wavefronts with a phase contribution from the object overlap. This means the returns will coherently sum producing a focused image of the object, but in an incorrect location, displaced in cross-range. Similarly, if the object is moving parallel to the radar (velocity purely in cross-range), the wavefronts with phase contributions will never have a single point where they all intersect. As a result the object will appear in the correct cross-range location, but smeared in the range dimension. In both cases the amount of displacement and smearing directly coincide with the relative radial velocity. Figure 2.8 illustrates the differences between parallel and perpendicular motion.

This means most moving objects in a SAR image will appear both smeared and displaced. The problem is further aggravated by the fact the target has four dimensions it operates in: range location, cross-range location, range velocity, and cross-range velocity. Combine this with the three dimensions measured by the SAR

platform, range location, cross-range location, and radial velocity, and suddenly a moving target can generate a given smear and displacement for any actual location given the correct radial velocity. If the target's motion parameters are unknown it becomes a time consuming and difficult process to correct for the motion and generate a focused image of the target in the correct physical location.

2.5 *Change Detection in SAR Imagery*

The first step in almost all currently used methods of correcting for motion in a SAR image is to isolate the returns of the moving objects from those of the static objects. For large changes over long periods of time, it is possible to implement non-coherent change detection on two generated images taken at different points in time. As non-coherent change detection is nearly identical to optical image change detection this will highlight the components of the images themselves that have changed and is useful for studying weather effects or construction progress [3, 6, 19, 22]. This is less useful for detecting smaller changes or isolating data from moving objects within the single scene. Since radar returns consist of complex data, coherent change detection makes use of both the phase as well as the magnitude information in an attempt to identify relatively small physical or temporal changes. Explained in the following subsections are two of the more popular coherent change detection methods available to SAR processing.

2.5.1 Interferometric SAR (IF-SAR). Similar to the non-coherent change detection methods there is a coherent method that is implemented on final generated images from two different collections. Interferometric SAR (IF-SAR) utilizes one radar both transmitting and receiving and a second calibrated radar only receiving. Both radar platforms target the same image scene but fly slightly different routes. This allows two images to be generated from the common transmitter, but with different phase centers. Utilizing the phase information, one complex image is multiplied by the conjugate of the second image, generating an IF-SAR image. If $f_m(x, y)$ is the

complex image generated by the transceiving radar and $f_b^*(x, y)$ is the conjugate of the complex image generated by the receiving radar, then the complex IF-SAR image, $f_i(x, y)$, as defined in [28] is

$$f_i(x, y) = f_m(x, y)f_b^*(x, y) \quad (2.20)$$

where x and y are rectangular spatial coordinates within the images.

This IF-SAR image contains information on the altitude of reflections in the scene based on the combined phase information. If two IF-SAR images of the same scene, taken at different times, are compared, the phase differences will reveal any changes in the scene. Rather than fly two separate aircraft it is possible to fly a single aircraft with a multi-channel radar system as long as the channels are sufficiently spaced to provide an adequate phase center displacement.

A second interferometric change detection method utilizes a single IF-SAR image in an attempt to identify moving objects within a scene. The IF-SAR image itself presents a continuous phase map of the targeted scene. Moving objects are displaced in location, however their phase information will be consistent with their true location. This means objects with a high radial velocity, with respect to the current radar location, will appear as discontinuities in the phase of the IF-SAR image [25]. Interferometry can identify subtle changes in the scene over time but still lacks the ability to isolate the returns of slow moving targets.

2.5.2 Displaced Phase Center Algorithm (DPCA). The displaced phase center algorithm (DPCA) is simple in concept but difficult to implement. A two or more channel radar flies a standard SAR collection path with the forward channel collecting data as described in Section 2.2.2. As the trailing channels arrive at the same physical location of the first channel, they also transmit a pulse, as shown in Figure 2.9. In this fashion it appears as if all channels of the radar collect data from exactly the same locations. If the channels all have a coherent phase source and identical beam

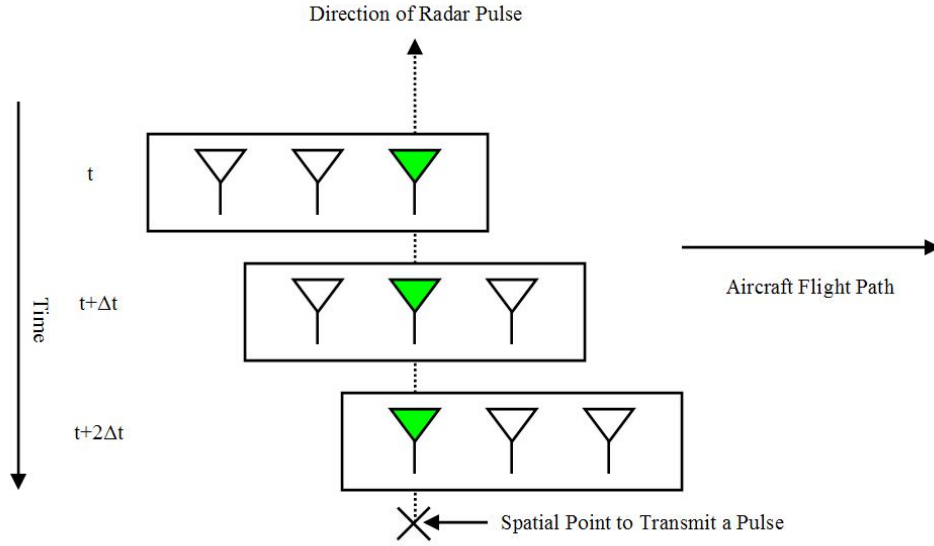


Figure 2.9: The displaced phase center algorithm relies on multiple radar channels sampling at a common point in space. A pulse is transmitted as each antenna, or channel, in this example three channel array arrives at the specific spatial point.

patterns, the collected data for all static items in the scene will be identical. Given it takes time for one channel to move to the sampling location, however, moving objects will be in a different location and provide a different phase contribution. By coherently subtracting one channel from another the static contributors will be canceled out and only the information from the moving objects will remain [4].

In ideal conditions, DPCA completely removes all static clutter, allowing slow moving targets to be clearly identified in the data. Unfortunately, DPCA is also highly susceptible to error and proves difficult to accurately implement. For instance, the channels must have identical beam patterns, must sample at exactly the same location in space, and transmit at exactly the same phase and frequency across the pulse. If any of these are not perfect, such as drift in the oscillator or aircraft flight path or imperfect switching in the radar sampling, the resultant image will have many streaks of static clutter contaminating the result [31]. If the error is high enough, it can become impossible to discern the actual moving targets from the error component.

2.6 *Shadow Detection*

Just as light sources create shadows behind objects in optical images, SAR images have shadows as well. A shadow in a SAR image represents a location where no energy is reflected back to the sensor. A shadow is not necessarily an obstruction of the transmitted energy by an object. Since no energy is reflected from the shadowed point, no phase information is available for that location. If the shadow was caused by an object, the shadow will appear in the correct location regardless if the object is moving or not. As mentioned previously, displacement of a moving object in a SAR image occurs due to the coherent summation of the returns as it moves through the collection. Without phase information, the shadow cannot have a distorted summation causing it to appear in the wrong location.

However, shadows of fast moving objects in SAR imagery will not appear as a typical dark spot in the image. As the object moves the locations obstructed will change, meaning at some points during the collection a shadowed area will reflect energy back to the radar and contribute a return. The total appearance of the shadow may only be a slight reduction in the energy level compared to the surrounding area, but it will not be physically displaced from the actual location of the moving object. This can be seen in a simulated DPCA image containing both static clutter and a moving target. The target will appear extracted from the clutter, but so will the shadow location, as shown in Figure 2.10. The contribution of the obstructed clutter will change over time, as the shadow passes over it, causing the energy from the clutter to not be completely canceled in the coherent subtraction.

If a target's shadow can be detected, the actual location of the target can be determined by projecting backward along the known illumination path (look angle). If the actual location is known, the target can be tracked or its radial velocity determined by comparing the apparent location to the known location. With the motion parameters and physical location known, it would be possible to refocus the target.

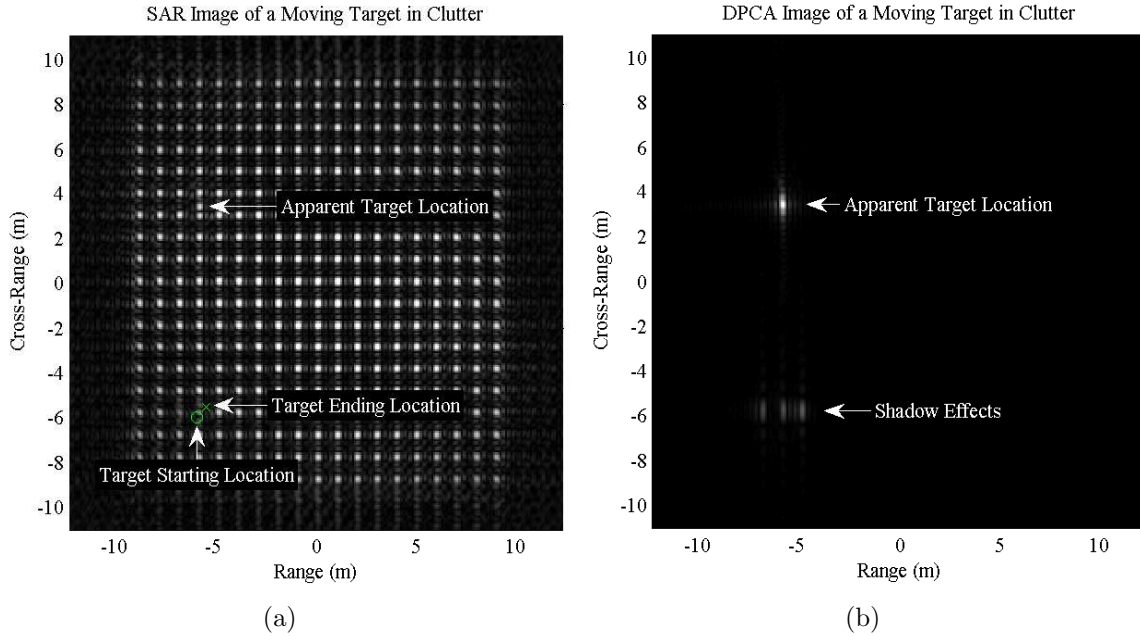


Figure 2.10: (a) The moving target point is taller than the surrounding clutter points so will cast a shadow across some of the clutter as it moves. (b) DPCA extracts the energy from the apparent moving target, but it also reveals changes in energy due to the target's shadow passing over the clutter. These shadow effects are much closer to the true location of the target.

Identifying both a target's return and its shadow location can possibly eliminate the problems associated with moving targets in a SAR scene.

2.6.1 Previous Work Exploiting Shadow Information in SAR Data. Most of the previous work regarding shadows in SAR data involved using shadow dimensions to extrapolate target information. In all cases the targets had already been located and the shadows were then isolated and used to either classify the target or extract information about the size and shape of the target. The authors in [2] and [9], for instance, either rely on humans to locate and delineate the shadow regions or assume shadows near a located target to be the region of lowest energy directly adjacent to the target. They use a "rattling rectangle" or "statistical snake" active contour delineation to define the boundary of the lower energy shadow region, and then use its extents to determine the maximum height and some limited shape information of

the target that cast that shadow. The authors of [5] and [26] similarly used shadow information to assist in target classification. However, in all cases the targets had already been identified by some other means and the shadows were used, after target location, as support to determine more information about the target.

2.6.2 Previous Work in Shadow Detection in SAR Data. Previous work in locating shadows in SAR data that exploits signal absence, or shadow, to locate and track moving targets was accomplished in [10]. In their work, the authors generate multiple images of a scene from sequentially collected data. These images are non-coherently summed to generate a single scene minimizing any contributions from moving objects and their shadows. Next, they create a ratio image by dividing the newly created reference image by one of the contributing SAR images. Any changes between the single image and the reference image will result in bright pixels. A spatial clustering algorithm is then run on this ratio image to determine sizes of shadows and produce possible targets at shadow locations of the correct size for the desired target dimensions.

The authors in [10] demonstrated the feasibility of their method if a large number of images from the collection are available to create the reference image. This insinuates any identification capability will require a large lead time to both collect enough information for the reference image *and* build the individual constituent images. This thesis proposes that target shadows can be identified using a single image thus creating an identification capability nearly as fast as the images can be generated.

Another method utilizing shadows in SAR imagery to locate targets is presented in [17]. The authors of [17] reason that the shadow of an object will be adjacent to the object itself, and similar to the research in this thesis, they look for this arrangement in the image to determine the locations of shadow casting objects. A sliding window across the image determines the means and standard deviations for both central group of pixels and the pixels immediately surrounding the central group. The ratio of the means of these two regions, by comparison to a threshold value, determines if the

central region is a “shadow” or a “target”. The locations of all shadows and targets are then compared and any of the proper size and arrangement (a target up-range and in close proximity to a shadow) are flagged as targets.

The method described in [17] requires down-sampling the image data because the sliding window is computationally intensive as it operates on a pixel by pixel basis. Additionally, [17] reports having difficulty locating targets in non-homogeneous regions where the shadow of the target itself can be obscured by surrounding shadows or the shadows of very large objects. This thesis proposes that not only can targets be located regardless of surrounding shadows from other objects, but the image can be operated on at the full resolution it was generated to.

2.6.3 Scene Statistics. The research described in this thesis relies heavily on estimation and detection theory. In order to utilize this basic set of theory, it is important to understand the statistics of the data. Shadows, as a lack of signal, will only contain system noise. Small targets and clutter will be affected by the amount of energy they reflect (signal), noise, and speckle. To identify the difference between a close grouping of low intensity targets and background clutter, such as an open field or road, it is important to understand speckle.

Speckle is defined by [20] as the noiselike characteristic of SAR images. Speckle is not noise in itself, only noiselike, and this fact is exploited later in Chapter III to assist in return characterization. The intensity of the speckle component, n , for a single look SAR image has an exponential probability distribution function (PDF) described in [20] as

$$P(n) = e^{-n}, \quad n \geq 0 \quad (2.21)$$

The intensity (or signal information) of a pixel in a SAR image is the product of the reflected energy due to the object’s radar cross section (RCS) value, σ , and the speckle intensity, n , defined in [20] as

$$I = \sigma n. \quad (2.22)$$

For targets with relatively constant RCS across large spatial areas, such as roads and open fields, the combination of RCS and speckle yields a multiplicative exponential PDF for the signal intensity [20] of

$$P(I) = \prod_{k=1}^N \frac{1}{\sigma} e^{-I_k/\sigma}, \quad I_k \geq 0 \forall k \quad (2.23)$$

where the target consists of N pixels and k is an index to each of the N pixels in the area of interest. This represents the joint PDF of independent and identically distributed intensity values for the N pixels.

A method of reducing speckle is to create a multi-look image. This is when multiple single images, typically two to four, from closely related angles are generated and averaged together. The effect of generating a multi-look image is that the PDF of the speckle changes from that of Equation 2.21 to a higher order Gamma distribution. The multi-look speckle PDF, defined in [20], is

$$P(n) = \frac{L^L n^{L-1}}{\Gamma(L)} e^{-Ln} \quad (2.24)$$

where L is the number of looks used to generate the multi-look image and $\Gamma(\cdot)$ is the Gamma function. As with single look images, when the RCS and the speckle components are combined per Equation 2.22, the resultant PDF has a similar Gamma distribution as that of the speckle alone. In fact, the averaging effect of multi-look image generation causes the image intensity PDFs of speckle, system noise, and constant RCS targets to all shift toward a Gamma distribution.

2.7 Summary

This chapter presented the principles of radar and extended the theory to SAR. It covered two common methods of SAR image generation and the problems encountered by moving objects in a SAR collection scene. Given a basic understanding of why the problems exist, a few methods used to mitigate those problems were discussed. Next, the reasons for working with signal absence rather than signal reflections were explained, as well as a review of the current work in this field. Utilizing the information presented in this chapter, Chapter III will detail a method for characterizing a SAR image and to identify target shadows based on these characterizations.

III. Research Methodology

This chapter presents the methodology and process defined in this research. The reasoning behind the requirement for scene characterization is presented, as well as a method of conducting said characterization. Utilizing the segment distributions from the characterization, a process of identifying shadows and the objects casting them is outlined and explained in a step by step description of the algorithm. Once the shadow-casting objects in the SAR imagery have been located, a method is introduced to filter the returns to include only objects of certain size.

The filtering process requires multiple parameters that directly relate to the type of object one is “looking” for. Since the parameters associated with each type of vehicle are initially unknown, a method of determining appropriate values is required. A heuristic is created as a “measure of goodness” of the detection rate associated with the filter. Finally, an optimization routine is selected to maximize the heuristic and determine the values of each of the filter parameters.

3.1 Shadow Detection

The shadow detection process is outlined in Figure 3.1. The following sections step through the process in detail, but a brief overview is presented here. First, an image, or set of images, is generated from SAR data collected for a scene of interest. As Section 2.6.3 covered, the exact methods used to generate the imagery will have an impact on its pixel intensity distributions. To eliminate any dependence on how the imagery was generated, the imagery must be characterized. This thesis assumes a human manually identifies each of the characterized regions in the image.

The purpose of the characterization is to identify the distribution parameters of shadows in the imagery. These parameters allow a shadow mask to be generated which highlights and separates “shadow” pixels from “non-shadow” pixels. Due to the properties of speckle, covered in Section 2.6.3, the shadow mask will be noisy. A low-pass speckle filter is applied to the mask to remove the noise and more clearly segment the shadow and non-shadow regions. Based on these regions and the orientation of

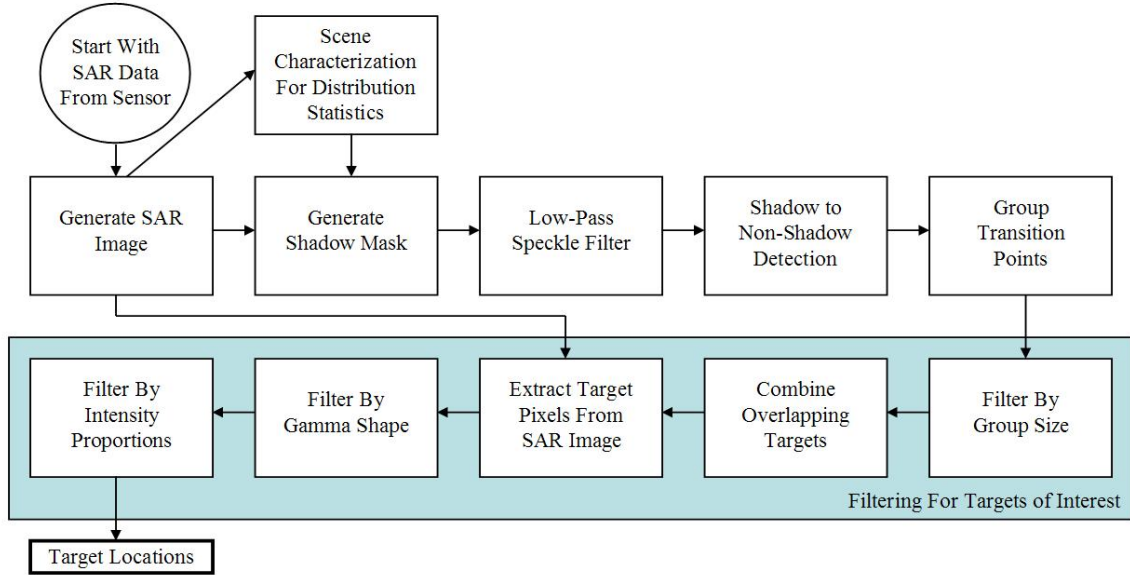


Figure 3.1: A block diagram of the target detection algorithm. Each block represents an operation in the process to detect targets based on shadow locations and pixel intensity distribution information.

the image relative to the radar, every pixel in the image is identified that represents a transition from shadow to non-shadow. Next, these transition pixels are spatially grouped, each representing a single shadow-casting object in the image.

Once every visible shadow-casting object has been identified, they are filtered to extract only the targets of interest. First, the groups are filtered based on size, including only those groups of the appropriate size given the desired target. Next, every overlapping target, based on group location and size, is combined to form a single target location. This creates a single target location, where before, portions of multiple targets of the desired size would have occupied the same physical location. Finally, the remaining targets are filtered based on the distribution parameters of their pixel intensities. Ideally, this results in the locations of every target of the desired type in the scene.

3.1.1 Scene Characterization For Distribution Statistics. The shadow detection process considers the distribution parameters of the associated with the pixel

intensities identified as shadows to be “*a priori*” knowledge as they require a human operator to segment and define a known region of shadow pixels once, while remaining constant from image to image of the same data set afterward. Although methods exist that attempt to characterize foreground (or target) from background (or non-target), these methods are not covered in this work as they are not applicable to the shadow detection problem. Shadows in SAR imagery are generally identified as an absence of energy with some contribution from noise. The contributors to the amount of noise are a combination of the sensor used for collection and the post-processing used to generate the image. Once a shadow is located it becomes important to know what created it. Objects with large areas of relatively constant RCS, such as trees, buildings, and fields, are generally of less interest than smaller objects like vehicles. To differentiate between the regions it is necessary to know their distribution parameters, which requires characterizing the image.

The foundation of the proposed process in this research is based on the assumption that different materials will have different distribution parameters for pixel intensity within the image. The objective of the characterization is to determine the parameters of these distributions. Before the parameters can be determined, the distribution type must first be discovered. For large areas with constant RCS, this will be either Gaussian or Gamma, depending on the level of post-processing performed on the image as covered in Section 2.6.3.

First, contiguous two-dimensional blocks of pixels known to be returns from the material to be characterized are extracted from the image and combined into one linear sample set. Then the type of distribution for this sample set needs to be determined. This can be accomplished by creating a histogram and “eye-balling” it or through automated “goodness of fit” methods such as the Anderson-Darling Test [1]. Once the type of distribution is known, a maximum likelihood estimation of the distribution parameters can be accomplished. So long as the post-processing procedures do not change from image to image, the characterizations from the first

image should be valid for any subsequent images generated from the same sensor collection.

3.1.2 Generate Shadow Mask. To extract shadow locations from the image, a binary distribution filter is created from the distribution parameters determined from the scene characterization process. The filter consists of a one-to-one image mask where the mask value for each pixel is the corresponding PDF model for that pixel’s intensity. The image mask is then converted to binary by determining the bounds that contain 95% of the distribution and setting all values inside those bounds to one, while all values outside those bounds are set to zero. This creates an image mask where all shadow, or shadow-like, pixels are clearly separated from non-shadows in the image. The 95% limit is chosen because some cutoff value is needed due to Gamma and Gaussian distributions being unbounded in the positive direction. If 100% of the intensities within the distribution were kept then all pixels in the image would be included. The value of 95% is selected because it is a common threshold for an upper limit of a confidence interval, and empirical testing of the 90% to 98% range suggests a value of 95% results in an acceptable balance of noise in the shadow and non-shadow regions. A more restrictive limitation could be used, but the resulting mask would contain more “holes” in regions of shadow where the pixel intensities of shadow regions fell outside the limits.

3.1.3 Low-pass Speckle Filter. Due to the high variability of speckle, the image mask is likely to contain a considerable amount of noise. This noise is represented by individual pixels throughout the image where speckle phenomenon causes the pixel intensity to cross into the shadow distribution. The noise in the speckle mask can be removed using standard noise filtering techniques. This thesis utilizes a low-pass Butterworth filter [15] to reduce the spatial noise. The mask is then returned to binary values by setting each mask pixel equal to or above one-half to one, and all pixels below one-half to zero. This will result in an image mask highlighting all

regions of the image that contain a majority of pixels in the intensity range of the shadow distribution.

Although there are many spatial filters and spatial filtering methods, the Butterworth filter is used in this thesis. The Butterworth filter is selected because it is a commonly used low-pass filter in image processing [23] and has a reasonable computational cost when filtering in the Fourier domain. The shape of the Butterworth filter is defined as

$$|H|^2 = \frac{1}{1 + C \frac{R^{2n}}{R_0^{2n}}} \quad (3.1)$$

where R is the distance from the center of the frequency spectrum, R_0 is the cutoff frequency, and n is the integer order of the filter. The parameter C is a constant that defines the value at the point $R = R_0$. For this thesis, C is set to one (as is often the case) defining the magnitude of that crossover point to be 50%. Since the filter operates based on a distance from the center of the spectrum, it can easily be adapted to two-dimensional filtering where the distance becomes a radius about the center of the two dimensional frequency spectrum. Figure 3.2 illustrates the magnitude response of a typical two-dimensional low-pass Butterworth filter.

3.1.4 Shadow-to-Non-Shadow Detection. A shadow-to-non-shadow detection is the identification of a location in the image mask where a shadow pixel is directly adjacent to a non-shadow pixel in the direction of the radar look angle. In any SAR collection the look angle is known for each pulse as it is directly related to the location of the scene center and the location of the sensor platform itself. This means the average look angle for each image generated is also known. The angle of illumination is the same as the look angle given the sensor itself is actively transmitting energy. Since this angle is known, the location of an object can be determined relative to the shadow it casts.




Figure 3.2: This first order Butterworth filter, with a normalized cutoff value of 0.5, has a gradual descent in its magnitude response, with the -3 dB line at the radial distance from the center equal to the cutoff value.

If the image, and thereby image mask, is aligned so that the sensor platform is at either the top or bottom of the image, the object/shadow relationship becomes one dimensional along the image's columns, as in Figure 3.3. The image mask is then iterated on a column by column basis. Starting at the pixel furthest from the sensor, the pixel is identified as either a shadow or non-shadow based on the binary value of the mask. If the pixel is labeled as shadow, a counter is incremented and the next pixel, moving closer to the sensor, is analyzed in a similar fashion. This continues until a non-shadow pixel is identified.

Since the resolution of the image is known, and the depression angle of a scene is also known from the height of the sensor platform and its distance relative to the scene center, the height of an object can be determined from the length of its shadow. Inversely, this means if the height of a desired target is known the length of its expected shadow can be estimated. Due to changes in terrain elevation affecting apparent shadow length, a range of possible shadow lengths could be compared to the value of the counter as a filter to exclude objects that are too tall or too short.




Figure 3.3: The SAR image is rotated or generated such that the average look angle of the image is perpendicular to the top and bottom edges of the image.

Once a non-shadow pixel is identified, the length of the associated shadow as determined by the value of the counter is compared to any potential length filters. If the shadow is of an appropriate length, the non-shadow pixel is flagged as a shadow-to-non-shadow transition. If the shadow is not of an appropriate length the pixel is not flagged. In either case the counter is reset to zero and processing continues pixel by pixel as before. When the end of a column is reached the counter is again reset to zero and the processing begins at the far end of the next column in the image mask.

By the end of this stage of the process, individual pixels have been identified that represent the transitions from shadows of a specified length to non-shadow.

3.1.5 Group Transition Points. The individual transition points offer very little information by themselves. To determine object size and to estimate the center of mass location, the points need to be connected to identify the edges of objects. This is accomplished by iterating a connected components process through the identified transition points. Starting with one corner of the image, the mask is iterated one pixel at a time, similar to the shadow-to-non-shadow detection step. When a transition pixel is encountered, two things occur. First, the pixel's group label is checked. If the transition does not belong to a group it is given the next available unique group identifier. Second, the eight surrounding adjacent pixels are checked for transitions. Any transitions identified are given the same group label as the center pixel.

Consider the example presented in Figure 3.4. Starting in Figure 3.4a, the pixel marked with a “?” is queried for a group label. In this case, a group label does not exist. In Figure 3.4b, the surrounding transition pixels are queried to see if any of them have a group label that could be applied to the first pixel, however none do. The first pixel originally under consideration is given the next available group label, which is 1, as shown in Figure 3.4c. In Figure 3.4d, the next pixel also has no group label, however one of the surrounding pixels does (the first pixel operated on.) The second pixel is then given the highest group label of all surrounding pixels, which is again 1, in this example. Figure 3.4e shows the first group completed and the next

directly reflects whether the identified transition line is either the center line of the object casting the shadow or the far edge of the object. The Moving and Stationary Target Acquisition and Recognition (MSTAR) data used in this thesis was taken at a 15 degree depression angle, meaning most objects of interest will have some amount of self shadowing and the center of the transition lines represent a good approximation to the location of the center of mass for the object.


At this point in the process, every object of a given height that casts a shadow has been identified. This includes buildings, trees, vehicles, signs, fences, and anything else that appears to cast a shadow matching the length of the given height restriction. This means not only have targets of potential interest been identified, but of potential false alarm sources as well. The next stage of the process filters out the false alarms from the desired targets.

3.2 Filtering For Targets of Interest


Now that most physical objects in the scene have been located, any number of identification or classification routines could be used to sift the results for those of interest. For the purpose of this thesis, vehicles are the desired targets. It should be possible to identify vehicles from non-vehicles by their distribution parameters.

3.2.1 Filter By Group Size and Combine Overlapping Targets. One of the simplest filters to implement is based on the size of the desired target. The transition lines that have been identified represent some fraction of the longest diagonal length of the object. By removing all transition lines that are too short or too long, a majority of the false alarms, such as buildings and signs, could be eliminated. However, due to self-shadowing, it is possible for a single vehicle to have multiple associated transition lines, as in the case depicted in Figure 3.6.

Since the size of the desired target is known, and the potential locations have been determined from the transition lines, bounding boxes can be created over the estimated target location at each transition line. The dimensions of the box are



(a)



(b)

Figure 3.5: (a) At a low depression angle self-shadowing can obscure a large portion of the overall target. (b) At a high depression angle more of the target can be imaged, but it is possible for the self-shadow to generate a second shadow-to-non-shadow transition line for the same target.




Figure 3.6: A target with two shadow transition lines of an appropriate width. In this case each line is the full width of the vehicle.

chosen based on the dimensions of the desired target. The size of the bounding box will affect the performance of the algorithm, as the statistics used later (Section 3.2.2) are based on the information contained within the boxes. The location of the box is placed at the center of the transition line. In the case of a single vehicle having multiple transition lines, this will result in overlapping boxes. To reduce false alarms and double positives any two overlapping bounding boxes are combined into one box located at the weighted average of their positions. The weight applied to each individual position is the length of the transition line. This causes the center of the resulting bounding box to be shifted toward the larger transition line in the hopes it is a better representation of the center of mass of the target. This process is then repeated until there are no remaining overlapping bounding boxes. In the case of tightly grouped vehicles, such as a parking lot, this might cause problems by combining two actual targets into only one detection. However, stationary tightly grouped vehicles should be clearly visible as such in the SAR imagery.

3.2.2 Filter By Distribution Properties. Now that potential targets have been reduced to objects of approximately the correct size, vehicles must be separated from non-vehicles. The RCS of a vehicle across its surface is extremely non-uniform [14]. This means the pixel intensity distribution of a vehicle will not likely be Gaussian or Gamma in nature. Based on the properties of RCS measurement and interaction presented in [14], it is assumed man-made objects with sharp edges and electromagnetically reflective surfaces, such as a vehicle, will tend to have only a few pixels with extremely bright intensity returns and much lower magnitudes across the majority of the remaining pixels.

Using the previously defined bounding boxes, segments of pixels are extracted from the SAR image as a representation of the potential target. Each box will contain a combination of either target and shadow or non-target and shadow. Given vehicles have sharp edges and are constructed of relatively flat surface reflectors, it is assumed that the majority of pixels for a vehicle should be much closer to zero than a potential target with a near uniform RCS, such as fields and level pavement [14,20]. For Gamma distributed targets, this can be measured by its “shape” parameter (see Section 2.6.3). Comparing the shape parameter of multiple distributions is equivalent to determining their relative skewness, assuming a constant scale between the distributions. This is because the equation for the skewness of a Gamma distribution is

$$\gamma_1 = \frac{2}{\sqrt{k}}, \quad k \geq 1 \quad (3.2)$$

where k is the shape parameter of the distribution. Figure 3.7 depicts constant scale Gamma distributions with varying shape parameters. There will be an unknown cutoff value for the shape parameter below which it is more likely a target than a false alarm. Determining the exact value for this cutoff will be explored in Section 3.3. Utilizing a maximum likelihood estimation to find Gamma distribution parameters on a non-Gamma distribution could be problematic. The shape parameter should provide a broad filtering mechanism, but further refinement will be required.




Figure 3.7: The scale parameter of each distribution is held constant at 1 while the shape parameter is varied for each plot.

The final filter implemented by this research to reduce false alarms again plays on the idea that vehicles will have only a small portion of their pixels at a high intensity, based on the RCS properties of complex shapes as presented in [14]. There can be no more than a certain percentage of pixels containing at least some fraction of the maximum intensity within the block. The method of identifying the values of these two percentages will be covered in Section 3.3. It is expected, from the above assumption regarding vehicle distributions, that vehicles will have the majority of their high intensity pixels restricted to a small percentage of the total.

The final result of the process is a set of locations within a SAR image representing shadow-casting objects that have been filtered to include only those of certain distribution properties. It is anticipated that these objects represent the targets of interest.

3.3 *Determining the Optimum Filter Parameters*

Three parameters for filtering based on distribution properties are listed as having unknown values in Section 3.2.2. The purpose of these filters is to maximize the number of desired target detections while minimizing the number of false alarms.

False alarms are defined as reported detections that are not the desired targets. The algorithm presented in this thesis operates at an object level, rather than pixel level. Each object is defined by a variable number of pixels, where the number of pixels is defined by the object itself. Essentially, this means the target detection problem is actually a two class classification problem, where any object is either “a target” or “not a target.”

Rather than developing a full contingency table for each instance of this classification problem, only the producers accuracy and consumers accuracy are calculated. These two values provide useful metrics as to how well the target filter is operating. In this two class problem, the producers accuracy is the same as its probability of detection. A classification routine’s probability of detection, P_D , is defined as

$$P_D = \frac{\text{Number of true targets detected}}{\text{Number of true targets available for detection}}. \quad (3.3)$$

This gives an indication to how well the classification filter correctly identifies targets. A second metric is needed to determine how often it incorrectly classifies non-targets as targets, generating false detections.

Due to how the algorithm in this thesis operates at an object level with varying object sizes across relatively large scenes, traditional methods of determining false alarm rates or probabilities of false alarm could result in misleadingly low numbers. The consumers accuracy, on the other hand, is a representation of the classification accuracy from the perspective of the user. It reveals what percentage of objects identified as targets are truly targets. However, since it is the number of non-targets identified as targets that is of interest, one minus the consumers accuracy is the metric used. This is denoted as the consumers error, CE , and is defined as

$$CE = 1 - \frac{\text{Number of true targets detected}}{(\text{Number of true targets detected}) + (\text{Number of false detections})}. \quad (3.4)$$

Rewritten to emphasize the false detections, this becomes

$$CE = \frac{\text{Number of false detections}}{\text{Number of reported detections}}. \quad (3.5)$$

In order to maximize the proposed process' P_D while minimizing its number of false detections, a heuristic is generated that relates the two. The heuristic needs to reward true detections and penalize false detections. The scaled probability heuristic, P_S , is therefore defined as

$$P_S = \frac{\text{Number of true targets detected}}{(\text{Number of false detections}) + (\text{Number of true targets available})}. \quad (3.6)$$

Similar to the probability of detection, assuming there are no false detections, P_S approaches one as the number of targets detected increases, and it approaches zero as the number of targets detected decreases. However, the number of false detections penalizes the heuristic such that its maximum value can only be obtained when there are no false alarms. As the number of false detections approaches infinity, P_S approaches zero. Examples of these cases are illustrated in Figure 3.8.

By running a known training set through the proposed process, P_S can be determined for any combination of the three unknown parameters. Determining the ideal parameter values then becomes an optimization problem attempting to maximize the scaled heuristic P_S . There are many optimization techniques that could be used for this such as genetic algorithms, tabu search, or stochastic tunneling. The method chosen for this optimization problem is simulated annealing.

Simulated annealing is chosen because it attempts to combine the efficiency of a hill climbing method with the completeness of a purely random walk [24]. Additionally, the function to be optimized is multi-variable and operates at the object level rather than pixel level, making it difficult to represent as a single equation. This means some other optimization methods, such as Neyman-Pearson, would be difficult




Figure 3.8: The heuristic, P_S , returns higher values for more target detections, but lower values as the number of false detections increases. (a) With no false detections, P_S behaves exactly as the P_D . (b) However, when false detections are added, the maximum possible value of the heuristic decreases, even if all targets are detected.

to implement and could return misleading results. Simulated annealing is a modification of the Metropolis algorithm which attempts to find the global minimum of a function by randomly picking a point in the function and comparing it to the current location. If the randomly selected result is better, it is automatically chosen. If it is not, then the point is chosen with some probability less than one, based on how “worse” the point is and the current “temperature” of the process. As time progresses the “temperature” decreases, in turn decreasing the likelihood a worse point will be chosen. The idea is that the large random component from the high “temperature” at the beginning will keep the process from selecting a local minimum, but as time progresses the function will eventually settle into the global minimum. Because simulated annealing utilizes a random component, the exact results will be different each time the method is applied to the same data set. For this reason multiple iterations are run on each training set and the best return determines the values to be used for the filter parameters.

```

Function SimulatedAnnealing(function, schedule)
    t = 0
    While T > 0 do
        T ← schedule(t)
        potential ← randomly selected from function with distance to best  $\propto T$ 
         $\Delta F = best - potential$ 
        If  $\Delta F > 0$ 
            best ← potential
        Else
            best ← potential with probability  $e^{\Delta F/T}$ 
        t ++
    Return best

```

Figure 3.9: Simulated annealing minimizes a function by replacing the current solution with a randomly selected point with a probability proportional to the supplied temperature schedule.

Simulated annealing can be represented by the pseudo-code shown in Figure 3.9 where *function* is the function to be optimized and *schedule* is a time schedule that gradually decreases the “temperature” until it finally reaches zero.

3.4 Summary

This chapter explained the necessity for statistically characterizing a SAR image as well as how to use the resulting characterizations to extract the locations of targets of interest. First, a method of characterization was described. Next, a process was outlined by which shadows are used to identify potential targets. Given it is desirable to isolate only a specific type of object that can cast a shadow, rather than all possible objects, a target filter was defined. The filter used the distribution statistics of the potential targets to separate vehicles from non-vehicles. Finally, a heuristic for measuring the accuracy of the filter was created and a method was selected for determining the filter parameter values that maximized the heuristic. Chapter IV presents the experiment methodology and the results at each stage of the process presented here.

IV. Experiment Methodology

This chapter presents the results of the target detection algorithm presented in Chapter III based on shadows in the Moving and Stationary Target Acquisition and Recognition (MSTAR) Synthetic Aperture Radar (SAR) dataset. The subset of the MSTAR data selected for processing is outlined, as well as the problems inherent to the methods used to generate the MSTAR imagery. The process described in Chapter III is begun by characterizing the intensity distributions of regions within the selected target chips and clutter scenes. The results of the characterization provide the foundations that allow the rest of the process to extract targets from the imagery.

The results for each step of the process are presented in sequential order, and establish whether or not potential target objects can be located from their shadows. Once potential targets are located from their shadows, a target/non-target discriminator is constructed based on the signal, or pixel intensities, of those potential targets. This discriminator acts as a filter to extract targets of interest from all of the available potential targets. The results of the discriminator are presented for each target and clutter scene type. Finally, examples of the process on sparse target scenes are provided.

4.1 Data Set

The data set used to evaluate the methodology specified in Chapter III is the publicly available data from the MSTAR program [18]. The MSTAR program was a joint venture between the Defense Research Advanced Projects Agency (DARPA) and the Air Force Research Laboratory (AFRL). The public data is a subset of a collection performed in September of 1995 over the Redstone Arsenal in Huntsville, Alabama. The sensor platform belonged to Sandia National Laboratory (SNL) and was operating in X-band at one foot resolution.

The MSTAR data set consists of pre-generated SAR imagery divided into clutter segments and isolated target chips. The clutter data was acquired in strip-map mode at a depression angle of approximately 15 degrees and divided into segments measuring

Table 4.1: Target Vehicle Dimensions

Vehicle	Height (m)	Width (m)	Length (m)
T-72	2.23	3.59	9.53
BMP-2	2.45	3.15	6.72
BTR-70	2.32	2.80	7.54

roughly 1,780 by 2,950 ft each. The target chips were acquired in spotlight mode and each measures roughly 128 ft square. Each chip represents a single target captured from a different look angle. Three vehicle types are represented in the target data. Focusing on only the target data captured at a 15 degree depression angle there are three Soviet T-72 tanks, three BMP-2 infantry fighting vehicles, and one BTR-70 armored personnel carrier. Each target vehicle is isolated and imaged from roughly 190 different incident angles. The dimensions of each vehicle will be important to identifying them as targets. Table 4.1 lists the volumetric dimensions of each of the three vehicle types.

4.1.1 Data Implementation. The methodology described in this section accounts for the differences in the methods of data collection and image generation between the target chips and clutter scenes. There are two factors relating to the images of interest that will determine if the differences in collection and image generation are too great to overcome. The first factor is resolution. This includes the resolution the scene is captured at and, more importantly, the resolution the image is generated to. The resolution determines the regional size a potential target and its shadow will cover. The larger the area the more defined the distribution, and the more likely the target can be detected. The MSTAR target chips and clutter scenes were generated to the same resolution. This means there is no scaling required to compare objects in both the clutter scenes and target chips, or to insert a target chip into any of the clutter scenes. The second factor is the amount of post processing performed on the image itself. Post processing techniques, such as speckle and noise reduction, can alter the distribution of pixel intensities within the image. As long as

these distributions are known, the proposed method can be implemented. Therefore, characterizing the distributions of selected regions in an image of the scene can remove dependence on specific post-processing, so long as each image to be analyzed has the same capture parameters and post-processing routines.

Unfortunately, the MSTAR clutter segments and target chips appear to have been generated with different parameters. Not only were the modes of capture different, strip-map for clutter vs spotlight for targets, the post processing was slightly different as well. The only indicator of a difference in post processing comes from analysis of the maximum intensities across the clutter segments and target chips. The target chips appear to have a much higher dynamic range than the clutter segments. However, the relative values between the two groups are comparable, and the sensor's magnitude calibration information, included in the header files, is constant between the target chips and clutter scenes. It will be shown later in this chapter that these apparent problems can be overcome by focusing on the magnitude-independent statistics when characterizing the scene.

4.1.2 Data Selection. As the MSTAR data is supplied with targets separated from the clutter, the two are kept separated while determining the effectiveness of the proposed process. Two clutter segments are used: HB06159 and HB06160. HB06159 represents a rural environment, containing mostly open fields with the occasional tree, fence line, and dirt road. HB06160 is a light urban setting consisting of a residential neighborhood with houses, roads, sidewalks, fences, and clusters of trees. Figure 4.1 shows the two clutter scenes chosen as they appear from the MSTAR data set before any processing from the algorithm presented in this research.

One representation of each of the three types of target vehicles included with the MSTAR data set is selected. The SN-132 set is selected for the T-72 tank. The SN-C71 set is selected for the BTR-70 personnel carrier (Note that this set is the only one available for this target type.) The SN-C21 set is chosen for the BMP-2. All three of the chosen target sets include chips of each vehicle imaged from 196 different


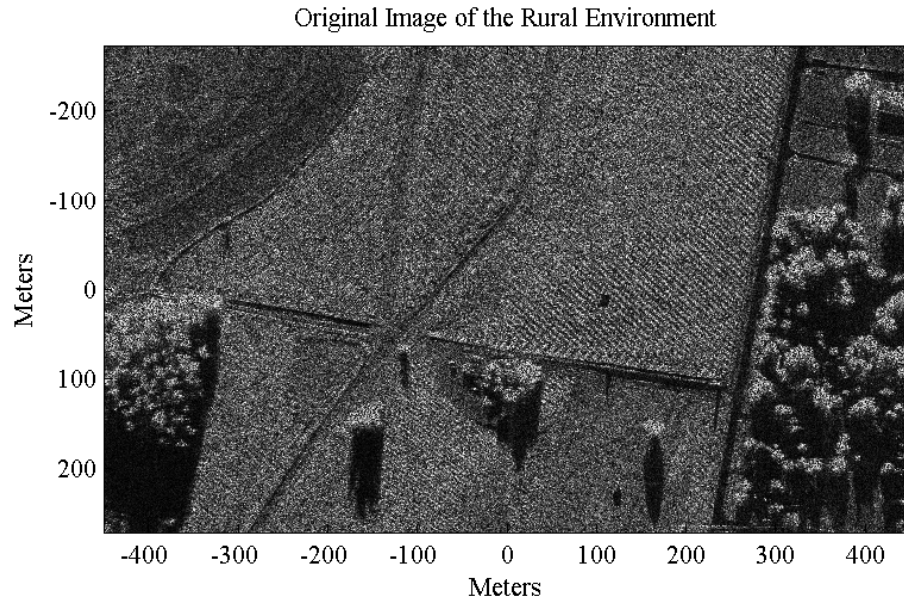


Figure 4.1: The presentation of the MSTAR clutter scenes used to evaluate the algorithms in this thesis. (a) HB06159 represents a rural environment. (b) HB06160 represents a light urban environment.

rotation angles. Figure 4.2 shows examples of the target chips for each of the three vehicle types.

4.2 Locating Targets From Shadows

The process of detecting targets based on their shadows is detailed in Figure 3.1, but a summary of each step is provided here. Then, the following sections will describe the results at each step. As described in the previous section, the imagery used in this thesis has already been generated. First, one of the clutter images is characterized for the distribution parameters of selected regions. Next, using the parameters of the known shadow region, a shadow mask is generated. This shadow mask separates shadow pixels from non-shadow pixels. A low-pass speckle filter is applied to the mask to remove the noise in the mask caused by the SAR speckle. Next, the shadow to non-shadow transition pixels are identified and grouped into clusters representing individual shadow casting objects.

The next segment of the process is to separate desired targets from clutter objects. Each group that is too large or too small to be of the desired target type is removed, and bounding boxes are created around the locations of the remaining groups. The size of the bounding boxes is determined by the dimensions of the desired target, and any that overlap are combined to form a single box. The statistics of the pixels that fall within these boxes are then determined and used to separate the desired targets of interest from the remaining background clutter.

For validation, and measuring the accuracy of locating targets based on their shadows, each vehicle is evaluated independently. Any filtering parameters of the process are set according to the expected size of the vehicle. Each target chip is then run through the process described in detail in Chapter III. Any targets returned from the process are considered “potential” targets. Assuming the targets in the chips are truly isolated, if multiple potentials are returned for a single target chip, the return containing the maximum total intensity is selected as the representation of the target, reducing the potential for false detections. The exact steps are outlined





Figure 4.2: Each of the three vehicles in these target chip examples stand out from the background, and their shadows are clearly defined. (a) A T-72 target chip. (b) A BTR-70 target chip. (c) A BMP-2 target chip.

in Section 4.3.2, but each target chip that returns at least one potential target is considered a positive detection. The same process is then run against the two clutter segments with the same parameters as used for each of the target chips and every potential target returned is considered a false detection (or false alarm).


4.2.1 Scene Characterization for Distribution Statistics. Scene characterization provides the basis for detecting shadows in the image. Every distinctly different region (such as grass, roads, trees, or shadows) in the image will have a different pixel distribution, as outlined in Section 3.1.1. The only distribution needed for this process is that of shadows cast by solid objects. However, the distributions themselves provide a sort of fingerprint as to the material they represent in the image. Figure 4.3 depicts the distribution histograms for various regions in the clutter image.

As expected, the histograms reveal that the pixel intensities of the regions with near constant RCS follow a Gamma distribution (e.g., Figures 4.3a and 4.3b). Using the histograms, each distribution’s parameters are determined to a 95% confidence interval, via maximum likelihood estimation. Figure 4.4 shows that each type of region’s distribution can be extremely different from the distributions of other regions. However, similar regions have similar distributions. The shadow distributions, for example, vary only slightly depending on shadow type. This is because the frequencies used in the MSTAR collection are not foliage penetrating and the shadows cast by solid objects are similar in nature to those cast by foliage. If the frequencies used were foliage penetrating, the distributions of the two shadow types would be dramatically different. Unfortunately, another region type with a similar distribution to solid object shadows is pavement. The distribution for pavement considerably overlaps both types of shadow distributions, and as such, it will be difficult to discern between the three regions in the imagery.


The difficulty in differentiating shadow from pavement is the most severe of the problems associated with the MSTAR data set, in terms of the objective of this thesis. The distributions depicted in Figure 4.4 were sampled from the clutter scenes. How-



(a)



(b)



(c)

Figure 4.3: (a) Pixel intensities for shadows cast by solid objects are dominated by noise in the system. The presence of a Gamma distribution here supports the assumption the MSTAR images had multi-look processing applied. (b) Grass lawn has a relatively uniform RCS leading to a clearly defined Gamma distribution. (c) Regions of trees have non-uniform RCS and so have a distribution that is neither Gamma nor Gaussian (it more closely matches a uniform distribution.)




Figure 4.4: The high energy return regions are distinctly different from the low energy shadow and pavement regions. However, shadows and pavement greatly overlap and are not separable.

ever, the target chips have a higher dynamic intensity range, and as such their shadow distributions are considerably different from those of the clutter scenes. Figure 4.5 reveals how drastically different the solid shadow distributions are between the clutter scenes and target chips. Since the characterization is done to the clutter scenes, and the targets were likely imaged resting on pavement, there will be an issue discerning the target shadow from the surrounding pavement in the target chips. Section 4.2.2 illustrates this issue, and Section 4.2.3 discusses how it affects the target location process.

4.2.2 Generate and Low-pass Speckle Filter the Shadow Mask. With the distribution parameters for solid shadows known, shadow masks are created for the clutter scenes and target chips to isolate “shadow” pixels from “non-shadow” pixels. Due to the characteristics of speckle in SAR imagery the resultant mask is extremely noisy, as shown in Figure 4.6.

A low-pass filter is needed to scrub the salt and pepper-like noise and clearly delineate the boundaries between shadow and non-shadow regions. This step is critical to the entire algorithm because every subsequent step is based on the resultant shadow




Figure 4.5: The shadows in the target chips have much lower average magnitudes than those in the clutter scenes. This makes the distributions of the two very different. The exact cause of the difference is unknown, but probably due to the different collection modes used (spot-light for targets and strip-map for clutter.)

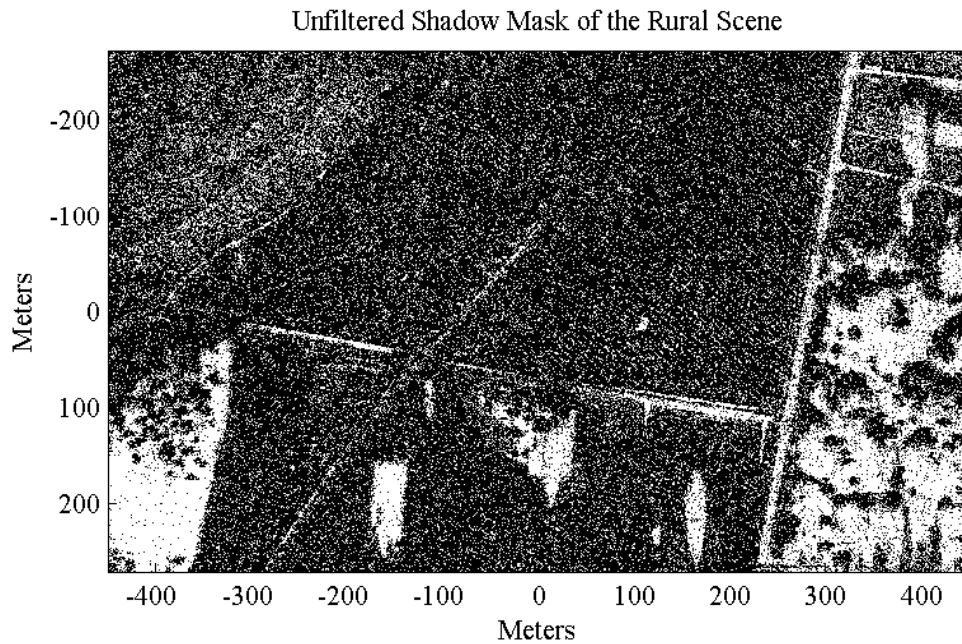


Figure 4.6: The initially generated shadow mask is extremely noisy due to the random aspects of speckle in SAR imagery. The white pixels are shadows and the black pixels are non-shadows.





Figure 4.7: The cut-off for each response plot is held constant at 50% of the normalized frequency while the order, n , is varied for each. Higher orders produce sharper drops in magnitude across the cut-off point.


mask. As described in detail in Section 3.1.3, a low-pass Butterworth filter is used for this thesis. However, other low-pass filters could produce equally useful results.

The Butterworth filter implemented is a fifth order filter with a normalized cutoff value of 0.1. The order of the filter acts to smooth the boundaries between the regions and the cutoff value determines the level of pixel variation considered noise to be removed. Higher order Butterworth filters provide sharper edge transitions, as shown in Figure 4.7. Therefore, a relatively high order Butterworth filter was chosen to create smooth edges between regions of shadow and non-shadow. Figure 4.8 contains the filtered shadow masks for both rural and urban clutter scenes. The large shadows from the clusters of trees and buildings are clearly visible. Unfortunately, so are the paved regions. Any region in the image that sufficiently absorbs or reflects the radar energy away from the sensor rather than scattering back to the sensor will be highlighted by the shadow mask.

Creating the filtered shadow mask for a target chip, shown in Figure 4.9, confirms the presence of the problem first mentioned in Section 4.2.1 where the targets will appear to be completely surrounded by shadow. Given the clutter scene characterization pointed out the overlapping distributions for shadows and pavement, this is the exact result expected for a vehicle driving down a road or resting in a parking



(a)



(b)

Figure 4.8: (a) The rural scene shows some noise in the open fields, but the mask is dominated by tree shadows and roads. (b) The urban scene contains far more shadows due to a higher density of objects that cast shadows, such as buildings and trees.




Figure 4.9: The T-72 tank stands out as the dark spot in the middle of the field of white. The shadow of the tank itself is obscured by the pavement with the similar intensity distribution surrounding it.

lot. This means the overlapping distributions do not truly represent a problem, just a worst case scenario where the shadow of the object is obscured in the shadow mask and no longer represents the object itself. The following section covers the method used to ensure the object itself is still detected in the almost overwhelming sea of shadow.

4.2.3 Shadow-to-Non-Shadow Detection and Group Transition Points. It is possible to locate every single shadow-to-non-shadow transition in the shadow mask. However, this would generate a large number of false detections due to any noise remaining after the shadow mask filtering step. To remove these noise-based false detections, the lengths of the shadows are taken into consideration as described in Section 3.1.4.

Compensating for the effects of possible elevation changes, the minimum shadow length allowed for each target vehicle type is 90% of the minimum length of the shadows cast on perfectly flat ground for each vehicle height (i.e., the shortest dimension

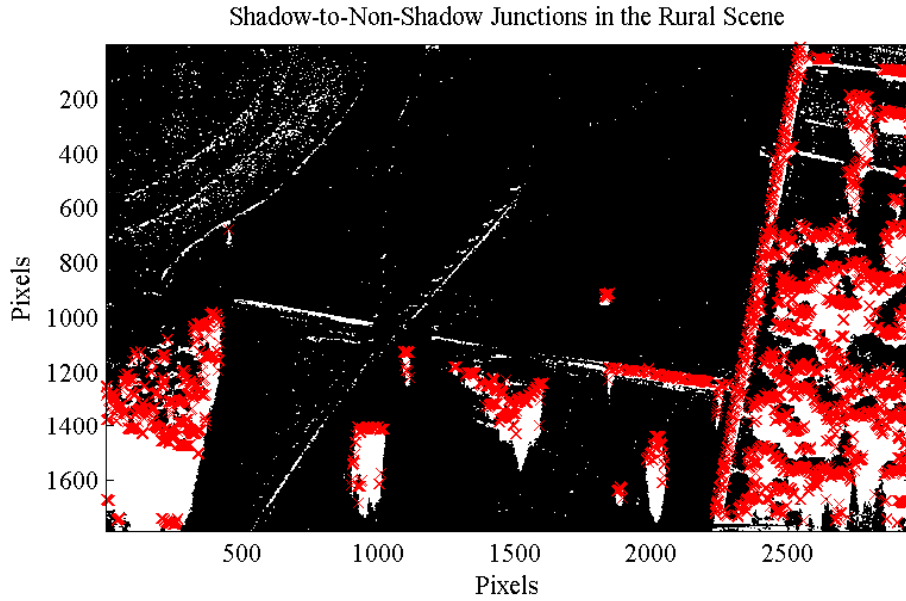


Figure 4.10: There are 6775 transitions corresponding to shadows of the appropriate length for the height of the T-72 tank.

of the vehicle is collinear with the look angle of the radar.) Unfortunately, since the shadows and pavement are both highlighted in the shadow mask, and the target chips are surrounded by “shadow-like” distributions, as shown in Section 4.2.2, there is no limit on the maximum shadow length allowed. This means even though the shadow is obscured by the pavement in the shadow mask, the transition from shadow to non-shadow at the edge of the target itself should still be identified.

The heights of each target vehicle are slightly different, meaning the number of shadow columns exceeding the minimum length in the clutter scenes could be different for each vehicle. Figure 4.10 shows every shadow-to-non-shadow transition in the rural scene for the T-72 tank highlighted by the red X’s. As described in Section 3.1.4, a shadow-to-non-shadow transition is a location in the image mask where a shadow pixel is directly adjacent to a non-shadow pixel in the direction of the radar look angle. Table 4.2 lists the number of transitions in each scene for each of the three target vehicle heights.

Table 4.2: Number of Transitions for Each Vehicle in Each Scene

Target Vehicle	Rural Scene	Urban Scene
T-72	6,775	24,779
BMP-2	6,231	23,026
BTR-70	6,775	24,779

Table 4.3: Number of Transition Groups for Each Vehicle in Each Scene

Target Vehicle	Rural Scene	Urban Scene
T-72	1,480	5,661
BMP-2	1,388	5,301
BTR-70	1,480	5,661

Once the individual transition pixels are identified, they are processed via connected components, resulting in pixel groupings that are, in turn, labeled with a unique group identifier. Table 4.3 lists the number of groups for each vehicle height in each scene type. As with the number of transitions, the number of groups for the T-72 and BTR-70 are identical because their heights comprise the same number of pixels due to the resolution of the image. The number of groups represents the maximum number of observable shadow-casting objects within the scene.

4.3 *Filter For Targets of Interest*

Well over 1,000 potential shadow casting objects are identified in each scene. Not only are actual objects such as trees, buildings, signs, and fences included, but transitions from roads to grass in the proper direction are included as well. From the definition of the data it is known that the thousands of detections in the clutter scenes are false detections. If a single target were dropped into either scene it would be obscured by the number of false detections and nothing would be gained. A filter is then needed to eliminate as many false detections as possible while retaining the actual target detections.

The method of filtering for targets of interest presented here is only one possible option. This filter is designed to identify known target vehicles using only distribution statistics, to overcome the different image generation methods of the data set. However, different data sets or different targets of interest could utilize drastically different filtration methods to remove the clutter. The following sections detail the results of the designed filter at each step in the process. The heuristic for the scaled probability of detection, P_S , as defined in Section 3.3, as well as the probability of detection and the consumers error (P_D and CE , respectively), are used in each filter step to determine that step's effectiveness.

4.3.1 Filter by Group Size. The sizes of each target vehicle are known and listed in Table 4.1. These lengths are converted to the appropriate number of pixels using the known resolution of the images. The shortest ideal shadow line for each vehicle will be the length of its shortest side. The longest ideal shadow line will be the diagonal length between opposite corners. In order to ensure all potential targets are identified, and take into account possible self-shadowing, the minimum and maximum lengths of the ideal shadow lines for each target vehicle are decreased and increased, respectively, by 50%. This 50% value creates a manageable range of possible shadow-to-non-shadow transition region lengths to bound the search space. The value of 50% may be increased to make for a more inclusive and time consuming search, or decreased to make for a more restrictive and comparatively faster search. Table 4.4 lists the minimum and maximum allowable shadow group length in pixels for each target vehicle type. However, to ensure targets completely surrounded by shadow are still detected, the maximum allowed shadow length for any target in this thesis is considered to be infinite rather than the values listed in Table 4.4.

Figure 4.11 reveals the effects of filtering the shadow groups by size in one of the T-72 target chips. Due to slight self-shadowing, there are shadow transition groups below the minimum length on either side of the main body. After filtering by the

Table 4.4: Shadow Group Limits, in Pixels, for Each Vehicle

Vehicle	Minimum Length	Maximum Length
T-72	5	51
BMP-2	5	39
BTR-70	4	41




Figure 4.11: (a) There are three separate shadow groups detected on this T-72 target chip (the yellow X's.) (b) After filtering by size, the two on either edge (caused by self-shadowing) are removed, leaving only the largest center of mass shadow transition group.

allowed length, those two groups are removed and only the shadow transition line of the main tank body and turret remains.

The sizing filter is run on all target chips and both clutter scenes. The number of shadow groups remaining, also referred to as detections, and the associated detection statistics are listed in Table 4.5. It is important to note that at this point some target chips are returning multiple detects per chip. This is handled in Section 4.3.2, but for now, since it is known that each chip should only contain a single target, any chip that returns at least one detection is considered only a single positive target detection.

Table 4.5: Size-Filtered Shadow Group Target Detection Rates

	Target Chips		Rural Scene			Urban Scene		
Vehicles	Detects	P_D	Detects	CE	P_S	Detects	CE	P_S
T-72	196	100%	458	70.0%	30.0%	1,886	90.6%	9.4%
BMP-2	196	100%	429	68.6%	31.4%	1,756	90.0%	10.0%
BTR-70	196	100%	561	74.1%	25.9%	2,314	92.2%	7.8%

Filtering by size removed over half of the false detections in all cases of the clutter scenes, but left every target detection intact. Even with the considerable improvement of results, the scaled probability of detection, P_S , reveals the output is still too cluttered with false detections to be of practical use.

4.3.2 Combine Overlapping Targets. The next step of the target filter is to use the known vehicle sizes to define bounding boxes about the transition lines that could contain the target and combine any of the boxes that overlap. Figure 3.6 provided an example of how a single target could have two shadow transition lines and Figure 4.12 reveals a T-72 target chip with this exact problem. The bounding boxes were created where each side is the average length of the minimum and maximum shadow length. This seemed the most logical, even though the target vehicles are rectangular in shape, because the orientation of the target is unknown and the desire is to capture the largest possible ratio of target to non-target pixels within the box. There are several approaches one may take to sizing the box. Ultimately, the size will affect the resulting performance as the statistics are based on the box formation (see Section 3.2.1).

After the bounding boxes are formed, overlapping bounding boxes are combined as described in Section 3.2.1, which should reduce the number of reported detections per target chip to one. However, to be certain of this, and in case a target chip contains something in addition to the desired target vehicle, the summed pixel intensity values in each returned bounding box are compared and the box containing the highest total

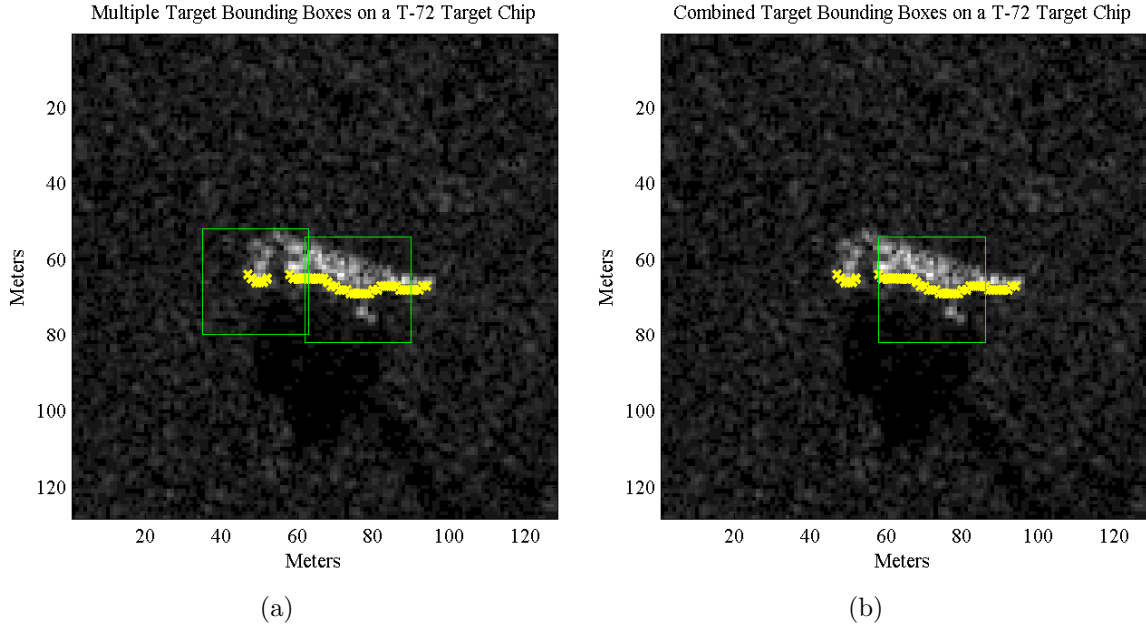


Figure 4.12: (a) Two potential target boxes from a single target will overlap. (b) Combining the two places a single box closer to the true center of the target.

intensity energy is selected as representing the target within the target chip. This guarantees every target chip with at least one detection returns only one detection, and that detection most likely contains the target vehicle and not an isolated support structure or radar noise. This is accomplished in order to produce adequate training data from the supplied target chips.

The combination process is run on each of the clutter scenes and the results and detection statistics are listed in Table 4.6. The detection rate and statistics for the target chips are omitted because the previous steps guaranteed any chip with at least one detection will still return a single detection. This means all 196 target chips for each vehicle contained one detection and retained the 100% probability of detection.

The number of false alarms have again been reduced by over 50% for the majority of different vehicle types and clutter scenes. However, the scaled probability of detection is still below 50%, especially in the urban scene. Therefore the filtering, to this point, is inadequate.


Table 4.6: Target Detection Rates After Combining Overlapping Bounding Boxes

	Rural Scene			Urban Scene		
Vehicles	Detects	CE	P_S	Detects	CE	P_S
T-72	211	51.84%	48.16%	866	81.54%	18.46%
BMP-2	242	55.25%	44.75%	984	83.39%	16.61%
BTR-70	275	58.39%	41.61%	1,072	84.54%	15.46%


4.3.3 Extract Target Intensity Statistics. Until now, each step in the filter has focused on the dimensions of the target and the detected shadow transition lines. This means the results currently contain both the target vehicles and any object in the clutter scene with a shadow transition line of approximately the same size. The next sections describe the pieces of the target filter that utilize the intensity distribution statistics of the individual detections.

The shadow mask was used to determine the locations and extents of the shadow transition lines. As was shown in Figures 4.11 and 4.12, the transition groups and potential target bounding boxes can be placed on the original image rather than the shadow mask used to determine their locations. Throughout the rest of this process, a target is described by the pixel intensities within the bounding box returned for each target chip. Similarly, a false alarm is described by the pixel intensities within any of the bounding boxes from each of the clutter scenes, as the clutter scenes contain no target vehicles and consist of only background clutter.


Histograms of the intensities for two of the target vehicles and a randomly chosen false alarm are shown in Figure 4.13. The shape of the vehicle distributions appear to be nearly identical to each other, yet narrower than that of the false alarm. This supports the assumption made in Section 3.2.2 that the shape parameter of the distribution, as well as the percentage of pixels below a certain threshold, can be used to separate the desired target vehicles from false detections.



(a)



(b)



(c)

Figure 4.13: (a) This example of a T-72 target retains the expected Gamma distribution. (b) This example of a BTR-70 target has an almost identical distribution to the T-72. (c) This example of a false alarm has a wider spread Gamma distribution than either of the target examples.

4.4 *Training the Filter*

Determining the optimum values of the three unknown filter parameters described in Section 3.2.2 (the shape of the intensity distribution, the percentage of the local maximum intensity to use as an upper bound, and the percentage of pixels with intensities beneath this bound) can be treated as a training problem. Each potential target is one of two classes, either a target vehicle or not a target vehicle. The goal of the filter is to correctly separate the potential targets into these two classes. As mentioned in Section 3.3, simulated annealing is chosen as the process to determine the optimum values the filter will use to make these separations. In essence, simulated annealing becomes a training process for a classifier comprised of the remaining steps of the target filter. Given simulated annealing finds the minimum of a function, such a function representing the accuracy of the filter is required.

4.4.1 Minimization Function. The scaled probability of detection, P_S (Eqn. 3.6), provides a scalar result, combining the detection rate and consumers error of the filter. Since the MSTAR data target chips and clutter scenes are kept separate throughout this thesis, the classifications of all available potential targets as either target or non-target are known. Therefore, the required function to be optimized is evaluated by simply iterating the remaining steps of the filter for any given three parameter values and determining P_S , the scaled probability of detection.

The target/non-target decision is based on the percentage of pixels that are above some pixel intensity value and the relationship between the expected and computed shape parameters of the pixel intensity distribution of the target chip. In the case of the pixel intensity value, an adaptive parameter is computed as a fraction of the chip's maximum intensity value. This fraction is defined as u . The image chip statistic of interest, N , is the fraction of pixels in the image chip above this adaptive parameter value and is defined as

$$N = \frac{1}{|X|} \sum_{i=1}^{|X|} 1(X_i \geq \max(X)u) \quad (4.1)$$

where X are the pixel intensities of the extracted image chip and $1(\cdot)$ is an indicator function that evaluates to one if its contents are true and zero if not true. The predicted upper limit to this fraction is the filter parameter n . Similarly, the Gamma shape parameter computed from X is denoted as K , while k is the filter parameter representing the upper limit of the allowed shape parameter values. This means a potential target is labeled a target if and only if $K < k$ and $N < n$ are both true for any selected set of the three parameters k , u , and n .

By classifying all available potential targets, and then comparing the predicted labels to the true labels of each, the scaled probability of detection can be determined per Eqn. 3.6. This, in turn, makes P_S a function of the three filter parameters k , u , and n . However, simulated annealing minimizes a given function, so the negative of the resultant $P_S(k, u, n)$ value is used, and the function to be optimized is complete.

4.4.2 Simulated Annealing Results. Since simulated annealing has a random component to it, each iteration can produce slightly different results. To ensure the best values possible are being returned, the simulated annealing optimization is repeated 100 times and only the parameter values corresponding to the best returned P_S are kept.

It was originally predicted in Section 3.2.2 that target vehicles would result in pixel intensity distributions with only a small number of intensities above a relatively high percentage of the local maximum intensity. However, the first run of the simulated annealing optimization for the BMP-2 produced the results shown in Table 4.7. These results contradict the original expectations of a large u and small n .

To verify the optimization function $P_S(k, u, n)$ worked as intended, and that simulated annealing was running correctly, P_S is plotted for varying values of u and n . Since only the pixel percentages u and n were of concern, k was held constant at

Table 4.7: Initial Simulated Annealing Results for the BMP-2 Target

Gamma Shape (k)	Maximum Intensity % (u)	Maximum Pixel % (n)
1.3941	5.4536	83.0450




Figure 4.14: The ridge showing the maximum possible values of P_S is close to the corner of the plot representing a low u and high n , opposite what was originally expected.

$+\infty$ and u and n were cycled through their range of possible values at 0.1% intervals. The resultant plot, using the BMP-2 data, is shown in Figure 4.14. This plot reveals two important facts. The first, and most important, is that the optimization function and optimization routine are both working as expected. This means part of the original assumption regarding the distribution of target vehicle pixel intensities, as they relate to other potential targets, was incorrect. The second, when compared to the results in Section 4.4.3, is that leaving k at $+\infty$ produces less accurate results than including a Gamma shape parameter value in the filter. This, in turn, means the other part of the original assumption was correct (the distribution shape can be used to help separate targets from non-targets.)

4.4.3 K-Fold Cross Validation. K -fold cross validation is a method of verifying the results of a classifier [7]. First, the members of known classes are each randomly segmented into K groups. The classifier is then trained using all but one group and the accuracy tested against the remaining group. This is repeated, cycling through which group is used for testing, until all groups have been used as the single testing group. Since the remaining portion of the filter operates as a two class classifier, determining if a potential target is either a target or non-target, K -fold cross validation is applicable. This results in a good estimate of the system performance.

The false alarms for the rural and urban scenes are combined for each target vehicle type, because high accuracy is desired regardless of the content of the background clutter. This results in each type of target vehicle retaining all 196 targets and having over 1,000 false detections, as defined in Section 4.3.3. Given the relatively low number of targets available for each vehicle, a 5-fold cross validation is used (a value of $K = 5$ or $K = 10$ is commonly used [7]). At this point it is unknown how similar the values of the three parameters will be between each vehicle type, so each vehicle is still kept separate and the 5-fold cross validation run for each. Tables 4.8, 4.9, and 4.10 list the results for the T-72, BMP-2, and BTR-70 respectively. The scaled probability of detection for all three vehicles is close to, or above, 80% with a probability of detection closer to 90%. This indicates the filter is accurately separating the targets from the non-targets with only a minimal loss in the number of detected actual targets.

In addition to the detection statistics, these tables list the values of the three filter parameters as returned by the simulated annealing optimization for that fold. These values are similar for all three vehicles, and almost identical between the BTR-70 and BMP-2. This suggests a broader filter could be created that detects military vehicles, as opposed to the specific vehicles as shown here.

Table 4.8: 5-Fold Cross Validation Results for the T-72 Data

Fold	P_S	P_D	CE	Shape (k)	Intensity (u)	Pixels (n)
1	96.88%	100.00%	3.13%	1.38	5.08%	82.44%
2	95.74%	95.74%	0.00%	1.35	5.08%	82.35%
3	95.35%	100.00%	4.65%	1.42	5.08%	82.39%
4	90.00%	90.00%	0.00%	1.37	5.04%	82.48%
5	90.38%	100.00%	9.62%	1.38	5.68%	79.87%
Mean (μ)	93.67%	97.15%	3.48%	1.38	5.19%	81.91%
SD (σ)	3.23%	4.40%	3.98%	0.03	0.27%	1.14%

Table 4.9: 5-Fold Cross Validation Results for the BMP-2 Data

Fold	P_S	P_D	CE	Shape (k)	Intensity (u)	Pixels (n)
1	82.00%	91.11%	10.87%	1.47	7.16%	71.72%
2	88.37%	95.00%	7.32%	1.50	5.74%	79.92%
3	80.00%	87.80%	10.00%	1.49	6.20%	76.60%
4	92.11%	94.59%	2.78%	1.47	5.02%	83.89%
5	86.11%	92.94%	8.82%	1.47	6.21%	76.40%
Mean (μ)	85.72%	92.49%	7.96%	1.48	6.07%	77.71%
SD (σ)	4.86%	3.03%	3.19%	0.02	0.78%	4.53%

Table 4.10: 5-Fold Cross Validation Results for the BTR-70 Data

Fold	P_S	P_D	CE	Shape (k)	Intensity (u)	Pixels (n)
1	76.32%	93.55%	19.44%	1.40	6.61%	78.18%
2	79.63%	84.31%	6.52%	1.40	5.42%	82.74%
3	87.10%	93.10%	6.90%	1.41	6.56%	78.26%
4	84.78%	86.67%	2.50%	1.47	5.51%	82.85%
5	81.82%	90.00%	10.00%	1.40	5.58%	82.88%
Mean (μ)	81.15%	89.53%	9.07%	1.42	5.94%	80.98%
SD (σ)	4.23%	4.02%	6.38%	0.03	0.60%	2.52%

4.5 *Mixed Targets*

Given the determined values of the three filter parameters are similar for all three vehicles it should be possible to create a slightly more inclusive filter that can isolate any of the three target vehicle types in a clutter scene. As a final verification that specific target types can be detected in SAR imagery based on the shadows in the image, a single target chip for each target vehicle type is selected at random and inserted into each of the clutter scenes. The values for the filter parameters are selected by choosing the average values from the 5-fold validation for each vehicle type that would eliminate the least number of potential targets. The filter parameters were originally determined to maximize P_S for a single specific vehicle type. By choosing inclusive values, the number of targets incorrectly labeled false detections should be minimized, but the number of false detections will also increase. This led to the assignments of $k = 1.48$ from the BTR-70 results, and $u = 5.19\%$ and $n = 81.91\%$ from the T-72 results. The dimension parameters of the earlier filter steps are then set to the height and longest possible shadow transition length of the T-72, and the shortest possible shadow transition length of the BTR-70. The entire shadow detection and target location process described in this thesis is then run on each combined image and the results are shown in Figure 4.15 where every potential target that the filter classifies as a target outlined in a green bounding box. Table 4.11 lists the results of each scene.

Not only do these images present the challenge of mixed target types, they also represent a sparse target environment. The previous filter detection statistics were generated with 196 targets of each vehicle. In each of these scenes there are only three targets total, yet the target detection process developed in this thesis detected all three in both clutter scenes. The low number of targets leads to a very low P_S for each image. However, the number of false alarms is low enough that a human operator could visually inspect each potential target to discern their validity, if needed. It is important to note that any potential target at the edge of a scene is only a fraction of the size of a potential target in the center of the scene which causes poor estimates





Figure 4.15: (a) All three target vehicles are identified in the rural scene. (b) Even though the number of false alarms is high compared to the number of targets, all three targets are identified in the urban scene, and there are few enough false alarms a human could easily validate them manually.

Table 4.11: Target Detection Results for Combined Target Scenes

Scene	Detects	False Alarms	P_S	P_D	CE	P_{SM}
Rural	8	5	37.50%	100%	62.50%	50.00%
Urban	19	16	15.79%	100%	84.21%	21.43%

for the distribution parameters, and thereby the distribution-based filter. The value of P_{SM} in the table is the scaled probability after the number of false detections are modified by ignoring any potential targets detected along the edges of each scene.

The previous research presented in [17] was also conducted on the MSTAR data set. Specifically, the clutter scene they used was a portion of the rural clutter scene used in this thesis. The research in this thesis utilizes only the target information in the two-class classification step. The results reported in [17] claim 14 false alarms due to clutter when only the target information is used. There are 4 false alarms in the same clutter region using the method and parameter values used to generate Table 4.11, which is a far better result than the 14 reported in [17]. The authors of [17] report only 2 false alarms when utilizing both target and shadow information in the classification. However, the trade-off is the inability to detect targets in non-homogeneous regions surrounded by other shadows, while the target detection algorithm in this thesis has no such problem.

As mentioned at the very start of the filter process, this filter was created based on the restrictions of the data set used and is meant to demonstrate that targets of interest can be located in SAR imagery based on the shadows in the scene. In this regard the process described in this thesis is a resounding success.

4.6 Summary

This chapter presented the benefits and limitations of characterizing the intensity distributions of SAR imagery. The characterization results can be used to isolate shadows within the SAR image, and thereby extract the locations of the objects casting those shadows. With the potential target locations isolated, it is shown that

specific targets of interest can be separated from the collection of shadow-casting objects based on the distribution characteristics of the desired target. This thesis demonstrates that it is possible to locate desired target classes within a SAR image utilizing the absence of signal, or shadows, in the scene.

V. Conclusions

This thesis introduces an algorithm that locates targets of interest in Synthetic Aperture Radar (SAR) imagery based on shadows in the image. Stationary objects in a SAR collection appear focused and in their actual geospatial locations. However, moving objects become unfocused and/or displace in the image from their actual locations. Traditionally, moving objects are refocused by altering the phase of the SAR samples in an attempt to compensate for the phase variation caused by target motion. These compensation techniques require knowledge of the motion parameters associated with the object. In the case of a non-cooperative target, these motion parameters are unknown and determining them requires cycling through the available search space for each of the parameters. Given objects have orientation relative to the radar and two dimensions each of velocity, acceleration, and change in acceleration, this can be computationally intensive and time consuming.

Section 2.2.1 explained that the displacement and blurring occur from the summation of the phase information in the SAR samples. Therefore, it stands to reason that shadows, which lack this phase information, should not be displaced from their true locations. This is the driving idea behind this thesis. If object shadows can be located within a SAR image, then the true locations of moving objects can be determined without having to rely on the object's motion parameters. This thesis demonstrates these ideas by first developing the shadow-detection of stationary targets.

5.1 *Research Contributions*

It is covered in Chapter II that very little publicly available work has been done, at the time of this thesis, in exploring the use of shadows in locating targets. This results in the need to analyze SAR imagery to determine if shadows contain enough information for use in target detection. The first research contribution of this thesis is demonstrating regions of near-constant RCS in SAR imagery can be identified

by the distribution parameters of their pixel intensities. This provides a method of consistently separating shadows, or shadow-like regions, from the rest of the image.

The second contribution is the demonstration that stationary objects can be located based on their shadows (or absence of signal) in the image, rather than the traditional method of utilizing the presence of signal. This establishes a foundation from which to pursue locating moving targets from the shadows they cast. Additionally, this provides a novel approach to locating stationary objects in an image, which can assist with target track initialization. By locating the stationary target, the track can be initialized immediately, rather than waiting for the target to begin moving, as is commonly done in track initiation processes.

5.2 *Future Research*

The algorithm in this thesis provides a number of options for areas of future research. The first is in improving the algorithm itself. As detailed in Section 3.1.3, the low-pass speckle filter used has a direct impact on the results of the algorithm. One step in the algorithm that could be explored is the development of a better low-pass speckle filter. Perhaps one that removes the speckle effects, but leaves exceptionally small shadows intact. Another option would be to explore image denoising techniques other than linear filtering. Ultimately, the objective would be to increase the accuracy of target detection, whatever the target or imagery used.

Another piece of the detection algorithm that could be expanded upon is the creation of the target bounding boxes. Section 3.2.1 describes the current method is based entirely on an average of the minimum and maximum dimensions of the desired target. Attempting to determine the orientation of the target based on its shadow, and thereby create a better “fitting” bounding box could have a dramatic impact on the accuracy of the algorithm.

The entire section of the algorithm dedicated to filtering the detection objects for targets of interest could also be replaced with a different set of routines. Rather than

sorting target from non-target, a filter could be designed that attempts to classify or identify all of the detected objects from a pool of known types. Target identification in SAR imagery has been the topic of a large amount of previous research that was not directly applicable to this thesis. However, future research could attempt to combine any of those routines with the algorithm presented here.

Future research could also attempt to include third party data sources into the algorithm. One example is if Geographic Information System (GIS) layers exist for an imaged scene, then masks based on those layers could be generated to eliminate false alarms. For instance, if the desired target is a vehicle, then false alarms coming from regions labeled as buildings in a GIS layer could be automatically eliminated. If the SAR sensor is flown in combination with a hyperspectral sensor these layers could be automatically generated based on, for example, the research in [8]. Another third party data source could be to utilize terrain elevation data from either an interferometric SAR collection or any topological data for the scene. The terrain elevations could be used to create a narrower, more accurate, range of shadow lengths for use in the length filtering step of the algorithm described in Section 3.1.4.

The final, and ultimately intended, area for future research is expanding the fundamentals of the shadow location algorithm in this thesis to the task of locating moving objects in SAR imagery. This thesis detects stationary objects by identifying shadows as regions of very little energy returns. The shadows of moving targets are dispersed over the region traversed by the object during the data collection. This means the faster a target is moving, the fainter the shadow will appear. Detecting these faint shadows will be more difficult than detecting the well defined shadows of stationary targets. However, it should be possible after modifying the algorithm as presented in this thesis. If a coherent change detection method is included with the sampled data, such as DPCA (explained in Section 2.5.2), both a static image and a phase change image could be generated. The change image highlights the displaced energy from moving objects as well as the shadow effects of the moving object. By analyzing the distribution parameters of regions in the static image where the change

image shows activity, it can be determined if the individual change is a shadow effect or a displaced and unfocused moving object. Shadows should have similar distribution parameters as the surrounding regions, except with a slightly lower average magnitude due to the blocked energy from the moving object. A shadow mask, similar to the one used for stationary objects in this thesis, can then be generated based on these regions of similar distributions but lower average magnitude. Using the locations of the shadow regions in the shadow mask and the look angle of the radar, it should be possible to then identify the location of the moving object in a similar manner to that described for stationary objects in Section 3.1.4. If refocusing the moving object is desired, in addition to simply knowing its actual location, then further research could attempt to reduce the search space for the motion parameters based on the shape and orientation of the associated shadow. Reducing the search space should decrease the computation time required to generate a focused image of the object.

Bibliography

1. Anderson, T. W. and D. A. Darling. "Asymptotic Theory of Certain Goodness of Fit Criteria Based on Stochastic Processes". *Annals of Mathematical Statistics*, 23:193–212, 1952.
2. Bennett, A.J. and D. Blacknell. "The extraction of building dimensions from high resolution SAR imagery". *Radar Conference, 2003. Proceedings of the International*, 182–187, Sept. 2003.
3. Bovolo, F. and L. Bruzzone. "A detail-preserving scale-driven approach to change detection in multitemporal SAR images". *Geoscience and Remote Sensing, IEEE Transactions on*, 43(12):2963–2972, Dec. 2005. ISSN 0196-2892.
4. Coe, D.J. and R.G. White. "Moving target detection in SAR imagery: experimental results". *Radar Conference, 1995., Record of the IEEE 1995 International*, 644–649, May 1995.
5. Cui, J., J. Gudnason, and M. Brookes. "Radar shadow and superresolution features for automatic recognition of MSTAR targets". *Radar Conference, 2005 IEEE International*, 534–539, May 2005.
6. Gamba, P., F. Dell'Acqua, and G. Lisini. "Change Detection of Multitemporal SAR Data in Urban Areas Combining Feature-Based and Pixel-Based Techniques". *Geoscience and Remote Sensing, IEEE Transactions on*, 44(10):2820–2827, Oct. 2006. ISSN 0196-2892.
7. Hastie, T., R. Tibshirani, and J. Friedman. *The Elements of Statistical Learning*. Springer Science+Business Media, Inc. New York, NY, USA, 2001.
8. Howard, T. E. *Abstracting GIS Shape Files From Hyperspectral Imagery*. Master's thesis, Air Force Institute of Technology, March 2009.
9. Jahangir, M., D. Blacknell, C.P. Moate, and R.D. Hill. "Extracting information from shadows in SAR imagery". *Machine Vision, 2007. ICMV 2007. International Conference on*, 107–112, Dec. 2007.
10. Jahangir, M. and C. P. Moate. "Utilising signal absence in SAR imagery for moving target detection". *Waveform Diversity and Design in Communications, Radar and Sonar, 2006. The Institution of Engineering and Technology Forum on*, 41–46, Nov. 2006.
11. Jakowatz, C.V., D.E. Wahl, P.H. Eichel, D.C. Ghiglia, and P.A. Thompson. *Spotlight-Mode Synthetic Aperture Radar: A Signal Processing Approach*. Kluwer Academic Publishers Norwell, MA, USA, 1996.
12. Kersten, P.R., R.W. Jansen, K. Luc, and T.L. Ainsworth. "Motion Analysis in SAR Images of Unfocused Objects Using TimeFrequency Methods". *Geoscience and Remote Sensing Letters, IEEE*, 4(4):527–531, Oct. 2007. ISSN 1545-598X.

13. Kirscht, M. "Detection and imaging of arbitrarily moving targets with single-channel SAR". *RADAR 2002*, 280–285, Oct. 2002.
14. Knott, E., J. Shaeffer, and M. Tuley. *Radar Cross Section*. SciTech Publishing, Inc. Raleigh, NC, USA, second edition, 2004.
15. Kovesi, P. D. "MATLAB and Octave Functions for Computer Vision and Image Processing". School of Computer Science & Software Engineering, The University of Western Australia. Available from: <<http://www.csse.uwa.edu.au/~pk/research/matlabfns/>>.
16. Levanon, N. and E. Mozeson. *Radar Signals*. John Wiley & Sons, Inc. Hoboken, NJ, USA, 2004.
17. Lombardo, P., M. Sciotti, and L.M. Kaplan. "SAR prescreening using both target and shadow information". *Radar Conference, 2001. Proceedings of the 2001 IEEE*, 147–152, 2001.
18. The Defense Research Advanced Projects Agency and the Air Force Research Laboratory. "Moving and Stationary Target Acquisition and Recognition (MSTAR) Public Release Data", 1995. Data available at <https://www.sdms.afrl.af.mil/datasets/mstar/index.php>.
19. Nonaka, T., T. Shibayama, H. Umakawa, and S. Uratsuka. "A comparison of the methods for the urban land cover change detection by high-resolution SAR data". *Geoscience and Remote Sensing Symposium, 2007. IGARSS 2007. IEEE International*, 3470–3473, July 2007.
20. Oliver, C. and S. Quegan. *Understanding Synthetic Aperture Radar Images*. SciTech Publishing, 2004.
21. Perry, R.P., R.C. DiPietro, and R.L. Fante. "SAR imaging of moving targets". *Aerospace and Electronic Systems, IEEE Transactions on*, 35(1):188–200, Jan 1999. ISSN 0018-9251.
22. Ranney, K. and M. Soumekh. "Adaptive change detection in coherent and non-coherent SAR imagery". *Radar Conference, 2005 IEEE International*, 195–200, May 2005.
23. Russ, J. C. *The Image Processing Handbook*. CRC Press Boca Raton, FL, USA, fifth edition, 2007.
24. Russell, S. J. and P. Norvig. *Artificial Intelligence: A Modern Approach*. Pearson Education, Inc., 2003.
25. Sanyal, P.K., D.M. Zasada, and R.P. Perry. "Tracking Moving Ground Targets from Airborne SAR via Keystoning and Multiple Phase Center Interferometry". *Radar Conference, 2007 IEEE*, 723–728, April 2007. ISSN 1097-5659.
26. Schumacher, R. and J. Schiller. "Non-cooperative target identification of battle-field targets - classification results based on SAR images". *Radar Conference, 2005 IEEE International*, 167–172, May 2005.

27. Skolnik, M. *Introduction to Radar Systems*. McGraw-Hill, 2001.
28. Soumekh, M. *Synthetic Aperture Radar Signal Processing with MATLAB Algorithms*. John Wiley & Sons, Inc. Hoboken, NJ, USA, 1999.
29. Soumekh, M. “Moving target detection and imaging using an X band along-track monopulse SAR”. *Aerospace and Electronic Systems, IEEE Transactions on*, 38(1):315–333, Jan 2002. ISSN 0018-9251.
30. Sun, H., G. Liu, H. Gu, and W. Su. “Application of the fractional Fourier transform to moving target detection in airborne SAR”. *Aerospace and Electronic Systems, IEEE Transactions on*, 38(4):1416–1424, Oct 2002. ISSN 0018-9251.
31. Zhang, H. and R. Yang. “Simulation of Multi-channel SAR Raw Data Based on Real Single Channel SAR Data”. *Radar, 2006. CIE '06. International Conference on*, 1–4, Oct. 2006.

REPORT DOCUMENTATION PAGE					<i>Form Approved</i> OMB No. 0704-0188	
The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.						
1. REPORT DATE (DD-MM-YYYY) 26-03-2009		2. REPORT TYPE Master's Thesis			3. DATES COVERED (From — To) Sept 2007 — Mar 2009	
4. TITLE AND SUBTITLE <div style="text-align: center;">Using Shadows to Detect Targets In Synthetic Aperture Radar Imagery</div>					5a. CONTRACT NUMBER 5b. GRANT NUMBER 5c. PROGRAM ELEMENT NUMBER 5d. PROJECT NUMBER 5e. TASK NUMBER 5f. WORK UNIT NUMBER 	
6. AUTHOR(S) Donnell, Brian P., Capt, USAF					8. PERFORMING ORGANIZATION REPORT NUMBER AFIT/GE/ENG/09-12	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Air Force Institute of Technology Graduate School of Engineering and Management (AFIT/EN) 2950 Hobson Way WPAFB OH 45433-7765					10. SPONSOR/MONITOR'S ACRONYM(S) AFRL/RYAS	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Air Force Research Laboratory/Sensors Directorate ATTN: Michael J. Minardi 2241 Avionics Circle WPAFB OH 45433 ((937) 904-9269; michael.minardi@wpafb.af.mil)					11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION / AVAILABILITY STATEMENT Approval for public release; distribution is unlimited.						
13. SUPPLEMENTARY NOTES						
14. ABSTRACT This thesis demonstrates the possibility of locating stationary targets of interest based on the locations of their shadows and the characteristics of pixel intensity distributions within synthetic aperture radar (SAR) imagery. Moving targets in SAR imagery are blurred and/or shifted in their apparent location. Shadows represent the absence of a signal reflection in a given location. The shadows of moving objects will always appear in their correct locations, providing clues as to the true locations of the moving objects. Three types of target vehicles are located in urban and rural clutter scenes using the publicly available Moving and Stationary Target Acquisition and Recognition (MSTAR) data set. The process described in this thesis allows stationary targets of interest to be located, which could assist in track initialization. The main purpose of this thesis is to build a foundation for locating moving targets based on their shadows. The distribution characterization method utilized demonstrates the feasibility of isolating distinct regions within SAR imagery and using the junctions between those regions to locate individual objects.						
15. SUBJECT TERMS synthetic aperture radar, radar, target detection, shadow						
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT		18. NUMBER OF PAGES	
a. REPORT	b. ABSTRACT	c. THIS PAGE	UU		96	
U	U	U	19a. NAME OF RESPONSIBLE PERSON Maj Michael J. Mendenhall, PhD (ENG)			
						19b. TELEPHONE NUMBER (include area code) (937) 255-3636; michael.mendenhall@afit.edu